<u>RipStream Predictive Harvest Effects Analysis</u>

Overview & ODF analysis approach

Part I: Predictive Shade & Temperature Model

1.1 Introduction the Predictive Shade & Temperature Model

1.2 RipStream study design background

1.3 PCW analysis

1.4 Effects analysis

1.5 Need for a more flexible approach – Bayesian modeling

1.6 Developing sub-models
- 1.6.1 Temperature
- 1.6.2 Shade

1.7 Combining the sub-models

1.8 Model evaluation

1.9 Model – specific assumptions

1.10 Model predictions

Part II: Harvest Simulation Approach

2.1 Vegetation data Use

2.2 Predictions
- 2.2.1 As-harvested
- 2.2.2 State Forests
- 2.2.3 Private Forests
- 2.2.4 Percent Basal Area
- 2.2.5 Harvest by distance from Stream

Appendix 1: Comparing 7DMM and 40-day max values

Appendix 2: Bayesian model parameterization

Appendix 3: State forest harvest simulation

Appendix 4: Private forest harvest simulation

Overview

In 1994 the Oregon Department of Forestry altered its Water Classification & Protection Rules (ODF 1994) to create the current suite of Oregon Forest Practices Act (FPA) riparian rules. In 1998 a MOU was signed between the Oregon Department of Environmental Quality (DEQ) and the Oregon Department of Forestry (ODF) that agreed to provide regulatory certainty for forest practices on non-federal land under the Clean Water Act so long as ODF monitoring verified that forest practices effectively protected water quality. ODF followed through on the agreement by enacting, in 2002, the Riparian Function and Stream Temperature project (RipStream). RipStream represented a highly controlled study of 33 sites in the Oregon Coast Range on state and privately-owned land. Data collection occurred generally for two years pre-harvest and five years post-harvest at all sites. Every site had a control reach upstream of the treatment reach. Data were collected repeatedly on stream temperatures, site vegetation, channel characteristics, and shade. Three peer-reviewed publications have come out of the study. Dent et al. (2008) provided an examination of site characteristics pre-harvest. The first post-harvest analysis, Groom et al. (2011a), asked whether the DEQ Protecting Cold Water criterion was met. It appeared that that PCW was not met on privately-owned forestland. The result warranted a closer examination of the data, resulting in a second post-harvest analysis (Groom et al. 2011b). This second post-harvest analysis had three main findings: 1) Streams on private lands appeared to be warmed by 0.7 ^{0}C while State forest streams did not change in temperature on average (0.0 ^{0}C); 2) Stream temperature changes were driven by changes in shade; 3) Changes in shade were largely related to changes in basal area. These results were presented to the Oregon Board of Forestry, the political body that oversees ODF and the administration of the FPA. The Board examined the findings and concluded that degradation of cold water had occurred on private lands. This finding triggered Section 527.714 of the Forest Practices Act, which initiates an examination of rule sufficiency and consideration of rule alteration. This paper presents the methods developed to produce recommendations to the Board of Forestry on alternative rule changes necessary to protect Cold Water.

The prediction method described here is intended for developing harvest prescriptions to the Board of Forestry. Although the sites were not randomly selected (see Groom et al. 2011b), virtually all available and suitable sites were used. The purpose of this effort is to model the temperature and shade responses to harvest at all 33 sites, and use the relationships found to predict temperature changes at those 33 sites given different harvest prescriptions. The overall temperature signal due to harvests across these sites will be interpreted as representative of harvests conducted across the geographical area of interest. The model is not intended for use beyond this effort; obtaining the necessary information, such as pre-harvest and upstream control temperature behavior, will generally not be possible in other settings.

Part I: Predictive Shade & Temperature Model

1.1 Introduction to the Predictive Shade & Temperature Model

The Predictive Shade & Temperature Model (PSTM) development depended on the earlier modeling efforts, yet it differs from them in some critical ways. The Protecting Cold Water analysis (PCW analysis; Groom et al. 2011a) produced findings that spurred subsequent analyses, yet it used temperature metrics that the others did not. The effects and magnitude analysis (effects analysis; Groom et al. 2011b) developed the shade and temperature models that were generally used by the PSTM described here.

1.2 RipStream background

1.2.1 Study overview

RipStream was conducted at 33 sites in the Oregon Coast Range. Sites were situated along first- to third-order streams on privately owned (18 sites) and state forest (15 sites) lands dominated by Douglas fir (*Pseudotsuga menzeisii*) and red alder (*Alnus rubra*). Forest stands were 50-70 years old and were fire- or harvest-regenerated. Openings were dominated by shrubs such as vine maple (*Acer circinatum*), stink currant (*Ribes bracteosum*), salmonberry (*Rubris spectabilis*), and devil's club (*Oplopanax horridus*).

Criteria to select sites included an ability to collect at least two years of pre-treatment and five years of post-treatment data at every site, the inclusion of sites harvested with Riparian Management Areas (RMAs) that meet current FPA and state forest standards, minimum treatment reach lengths of 1000 feet, and assurance that the upstream "control" reach would remain unharvested for the duration of the study. Streams needed to qualify under the FPA as "Small" or "Medium" (mean annual streamflow < 2 cfs or between 2 and 10 cfs, respectively), and streams needed to be free of recent impacts from debris torrents and active beaver ponds. We obtained sites by requesting that industrial private and state forest managers in the Oregon Coast Range provide ODF with a list of stream reaches that met the criteria and would be harvested no sooner than 2004. An initial list of 130 stream sites was reduced to 36 sites that met study design constraints. Three more were subsequently dropped due to changes in harvest plans. While there was an initial attempt to exclude sites with beaver activity, a beaver dam ponded 722 ft of the 3,806 ft treatment reach for site 7801 during the first and second post-harvest years.

Each site has a control reach immediately upstream of a treatment reach. The control reaches were continuously forested to a perpendicular slope distance of at least 200 feet from the average annual high water level. Reach lengths varied from 450 ft to 6000 ft with means of 905 ft and 2244 ft for the control and treatment reaches, respectively (Dent et al. 2008).

1.2.2 Treatments

Forest Practices Act On Private Sites: Eighteen sites were established on private forest streams. Sites were harvested following contemporary FPA standards which require riparian buffers along fish-bearing streams to protect stream temperature, provide future large wood for streams, and retain other ecological services (Oregon Department of Forestry 2007). Measured as a slope distance, the RMAs are 50 and 70 ft wide around small and medium fish-bearing streams, respectively. Both small and medium streams have a 20-ft no-cut zone immediately adjacent to the stream. Harvesting is allowed in the remaining RMA to a minimum basal area of 40 (small streams) and 120 (medium streams) ft^2/1000 ft. See section 2.6 for more details.

Oregon State Northwest Forest Management Plan (NWFMP) on State Sites: In Oregon, lands administered directly by Oregon Department of Forestry (state forests) are managed under FMPs for multiple resource objectives (e.g., recreation, wildlife) in addition to timber production and require riparian protections that exceed FPA minimum values. Fifteen sites were established on state forest lands. All but two RipStream state forest sites had RMAs managed according to the Northwest Forest Management Plan; two sites were managed according to the Elliott Forest Management Plan which has identical riparian management strategies. We therefore refer to the management of all 15 state sites as NWFMP. Measured as a horizontal distance from the stream's edge, state RMAs are 170 ft wide for all fish-bearing streams with a 25-ft no cut zone. Limited harvest is allowed between 25 ft and 100 ft of the stream only to create mature forest conditions. Further specifications are provided in section 2.7.

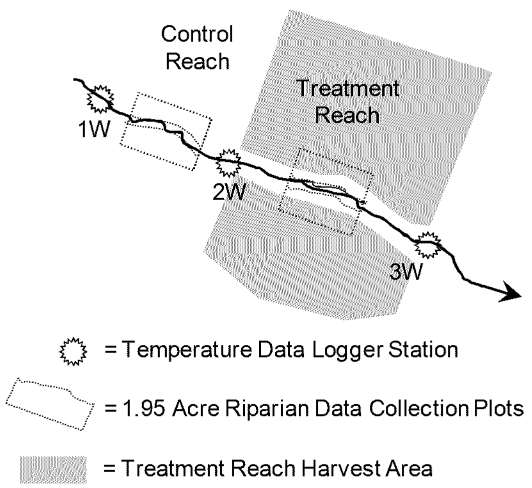

Figure 1. RipStream site layout. Control and treatment reaches are defined by the position of Station 2 and 3 temperature data logger stations that coincide with harvest boundaries. Two riparian data collection plots are situated midway along both the control and treatment reaches.

1.2.3. Data Collection

Optic Stowaway Temp and HOBO Water Temp Pro data loggers (Onset Computer Corporation, Bourne, Massachusetts) were annually deployed at four stations beginning in 2002 or 2003 (Figure 1). In the PSTM analysis we consider only Stations 1-3, not 4. 1W is at the upstream end of the control reach, 2W is located at the downstream end of the control reach and the upstream end of the treatment reach (i.e. shared logger) and 3W is at the downstream end of the treatment reach. Temperature loggers were deployed in shaded locations where stream flow was relatively constant, with reliable summer depth, and a well-mixed water column. Logger accuracy was checked prior to installation, in the field, and following retrieval with National Institute for Standards and Technology-calibrated digital thermometers according to the stream temperature protocol described by the Oregon Watershed Enhancement Board (1999). Although both types of loggers we used are listed as ±0.2 °C accuracy, we found that for over 500 pre- and post- deployment assessments in only two instances did loggers register errors of >0.1°C. Daily temperatures that exhibited increases in diel fluctuation and increases and decreases in daily maximum and minimum temperatures that were not reflected in other probes during the same year or at the same location during other years were interpreted as influenced by air temperature and excluded from the analysis.

Stream channel data were collected at 200 ft intervals within each reach. Data included wetted width, bankfull width, thalweg depth, and stream gradient collected according to the protocol described by Kauffmann and Robison (1998). Stream shade was quantified at these intervals using a self-leveling fisheye lens digital camera (Valverde and Silvertown, 1997). Shade values were measured once pre-harvest and once post-harvest. Fish-eye photographs were taken in the middle of the stream, 1 m above the water level, and oriented due north. Shade values were calculated from the photographs using HemiView™ 2.1 software (Delta-T Devices, Cambridge, UK) as one minus the June 30 Global Site Factor (1 -GSF). The GSF is the proportion of both direct and diffuse energy under a plant canopy relative to the available direct and diffuse energy for the given site's latitude/longitude. Shade and gradient values were averaged for each reach.

Vegetation data were collected in four 500 by 170 ft plots on both sides of a study stream in the treatment and control reach (Figure 1). Plots were centered midway along each reach and contoured according to stream curvature. Vegetation plot data describe understory, overstory, downed wood, blowdown, and snag characteristics. The original purpose of including extensive vegetation plot data collection was to assess large wood recruitment, shade, and riparian structure following timber harvest. For this analysis we use a portion of the available riparian structure data (e.g., blow down, tree heights, basal area, species). Within each plot all living trees with a diameter at breast height (DBH) > 6 inches were tallied by species; each tree's distance to the stream was additionally recorded. Height, live crown ratio, and crown class (e.g., dominant, co-dominant, intermediate, overtopped) were additionally measured for 20% of the trees. Figure 2 depicts the plot layout that determined data collection. Tree data were recorded in 100 x 170 ft "lines". In the middle of each line was a transect, along which riparian downed wood was tallied. Each of the 1/100th acre subplots (circles) within each line were situated at 25 ft increments beginning at the stream's edge. Within each of a plot's 30 subplots, contractors recorded landform, hillslope (measured towards the stream), shrub composition and percent cover, and seedling data.

Data were collected in all four vegetation plots per site pre-harvest and re-measured in harvested treatment plot or plots (if one or both stream sides were harvested, respectively) post-harvest. Blowdown was quantified in all plots post-harvest.

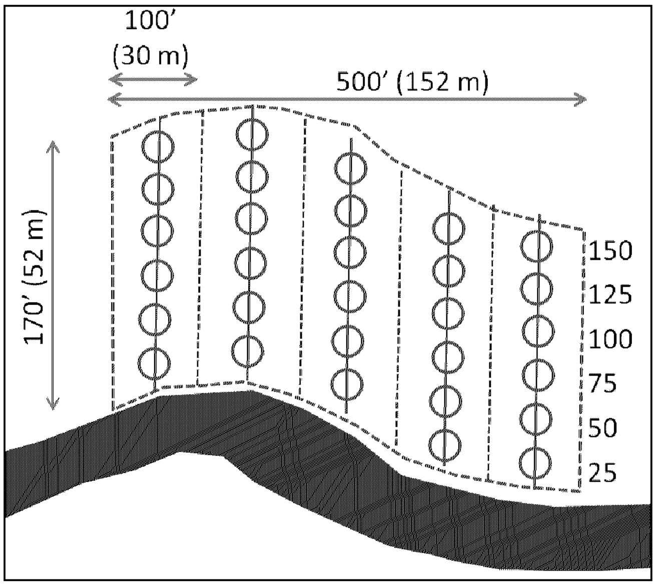

Figure 2. Data collection schematic for a single vegetation plot. Plots contoured stream channels. Within each plot are five 100 ft x 170 ft lines. Each line has a transect, and along each transect at 25 ft intervals are 1/100th acre sub-plots.

The PCW, effects, and PSTM analyses are limited to all pre-harvest data and data from the first and second post-harvest years. Data collection included hourly stream temperatures (collected annually between July 1 and September 15) and both channel data (for a complete list see Dent et al., 2008) and riparian vegetation data (overstory and understory) during the first years pre-harvest and post-harvest.

1.3 PCW analysis

Our publication on the PCW analysis (Groom et al. 2011a) presents an unusually complex analysis. The complexity stems from our goal of answering a regulatory question using a statistically valid analysis approach constrained by the structure of the regulation. Specifically, we wanted to know if stream temperatures observed in RipStream sites indicated whether timber harvest was increasing stream temperatures by more than 0.3 ^{o}C.

The Oregon Department of Environmental Quality water quality rules include several stream temperature criteria. Among them is the Protecting Cold Water (PCW) criterion that represents a federally-required antidegradation water quality rule component. The PCW applies to "cold" streams with temperatures below specific temperature thresholds

(DEQ, 2004). Anthropogenic activities such as timber harvest are not permitted to increase stream temperature by more than 0.3 ^{o}C above its ambient temperature.

The criterion is fairly straightforward to assess at a point source, such as below the outlet of a pipe. Samples would be taken immediately above the outlet and below the mixing zone past the outlet. However, forest harvests affect stream temperatures as a non-point source process. The PCW temperature metric of interest is the 7DMM, or seven-day mean of daily maximum temperatures. This value is quantified as a moving mean across a season. That is, the first 7DMM represents day 1 through 7, while the next 7DMM will represent days 2 through 8. The Protecting Cold Water criterion indicates that if any 7DMM is elevated by more than 0.3 ^{o}C above a baseline condition then the criterion is not met.

We performed an analysis of the PCW (Groom et al. 2011), reducing our hourly data to daily summaries, and then to 7DMM temperatures. To determine if any single 7DMM is out of compliance required us to establish a reference condition and then ask if any 7DMM during the test condition exceeded the expected amount by more than its prediction interval plus 0.3 C. For the detailed methods of the analysis, please see Groom et al. (2011a).We restricted the parameterization of the PCW analysis to adhere to rule language. The rules do not consider factors such as the length of the harvested reach, year effects, gradient, or other pertinent variables. Therefore our analysis similarly excluded these potential factors as well. In addition, the analysis examined every summer 7DMM datum relative to expected values. As a result, the analysis provided an answer to our analysis question (we saw an elevation of temperatures in private sites above the 0.3 ^{0}C threshold), but the analysis itself was convoluted enough that using it to obtain temperature change magnitudes would be meaningless or at best provide suspect results.

The results of the PCW analysis were presented to the Oregon Board of Forestry several times between 2009 and 2011. At the January 2012 meeting the Board of Forestry determined that the RipStream findings indicated forest practices contributed to the degradation of a natural resource, cold water. This finding triggered an FPA riparian rules analysis that potentially leads to a change in the riparian rules. This decision led ultimately to the development of the analysis approach presented here. However, the modeling effort and results of the PCW analysis are not incorporated in the PSTM due to factors discussed above. Instead, the PCW findings led to a second analysis, the effects analysis, which examined the magnitude of temperature increase at sites as well as important site variables that were related to observed temperature change. The effects analysis serves as the basis for the PSTM.

1.4 Effects analysis

In 2011 we published a second manuscript (Groom et al. 2011b) that delved into site variables and how they related to observed temperature change. This “effects analysis” abandoned the constraints of the PCW analysis. The effects analysis examines the contributions of different variables at explaining observed temperature changes, including treatment reach length, average treatment reach shade, elevation, average treatment reach gradient, east-west deviation in degrees of the treatment reach valley (a simplified version of valley azimuth), change in control reach temperature, and watershed area calculated at 3W. We also included state and private ownership and the harvest status or whether a temperature measurement occurred during a pre-harvest or post-harvest year. Combinations of these variables were used in a suite of mixed linear regression models. We determined that mixed models were called for, as such models allow site intercept and the parameter value associated with the control reach temperature change to differ by site, accommodating random site-specific characteristics. Using the same mixed parameter in the models we compared the performance of 18 a priori temperature models (Groom et al. 2011b Table 1).

We examined model performance for four temperature change metrics. We summarized hourly stream temperature data to provide daily maximum, mean, minimum, and fluctuation (maximum-minimum) values for each station. We were interested in detecting changes in stream temperature due to site factors including harvest. We therefore defined the response variable as the daily difference between treatment reach 2W and 3W. To reduce analysis complexity we computed the average of this difference over a forty day period for each year (July 23 to August 15). This represents the time frame when we had the greatest number of functional loggers recording temperatures during a central portion of the summer months when maximum temperatures are observed in the Oregon Coast Range. We compared the 18 *a priori* temperature models against these four temperature metrics.

The PCW analysis described in section 1.3 focused on maximum daily temperatures, averaged across seven day periods. Although the effects analysis examined a suite of temperature metrics, our current effort is focused on predicting effects of harvest on daily maximum temperatures as the DEQ temperature criteria focus on these quantities. Fortunately, the effects analysis' metric for maximum daily temperatures averaged over a 40-day period (40-day max) is virtually identical to the mean of the 7DMM values taken for the same time period (Appendix 1). Therefore, we interpret the 40-day max values as a suitable substitute for an average response of 7DMM values.

The analysis focuses on an average response of the 7DMM instead of individual values out of a need for model simplicity. Streams temperatures differ in how they change temporally and spatially. Had we used 7DMM values we would have needed to include some modeling component of 7DMM values to account for site- or even year-specific trends. The model would have also needed to incorporate autocorrelation and moving average corrections for these data. This level of complexity was avoided by using the average. A consequence, however, is that the 40-day max values do not capture the highest (or lowest) 7DMM values at sites.

The 40-day max values for all years pre-harvest and the first two years post-harvest at all sites were best explained by variables for the change in control-reach temperature (Control), treatment reach length (TRLength), the mean of treatment reach shade measurements (Shade), and the first quartile of stream gradient measurements (Grad1Q).

Control values were the 40-day max temperature change for a given year between the 1W and 2W probes. The values were normalized by control reach length. TRLength was determined from field measurements of the distance between 2W and 3W. Shade was measured as described in section 1.2.3. Gradient values appeared generally skewed by high-gradient values, so we conducted the analysis on the first quartile (lower readings) of channel gradient, relating to slower stream flow rates.

The formulation for the best-supported statistical model was:

$$[1]\ \Delta T_{3-2ijk} = \alpha_0 + \alpha_j + (\beta_1 \Delta TControl_{2-1i} + \beta_j \Delta TControl_{2-1i}) + \beta_2 TRLength_j + \beta_3 Shade_k + \beta_4 Grad1Q_j$$

Subscripts indicate year data *i*, site *j*, and pre- or post-harvest status *k*. The model includes mixed-effects parameters for the linear model's intercept (α_j) and the slope value for Control temperatures (β_j). These values allow a different intercept and Control value for each site, assisting in accounting for the lack of independence between 40-day max observations at a single site. The modeling effort determined that providing the structure of a mixed model was advantageous over an intercept-only mixed model or a standard linear regression.

Pre-harvest shade was constant enough between values of 80% and 95% that few variables could account for between-site differences. We therefore created linear models to describe only post-harvest shade. We constructed the models as weighted linear regression due to the different number of shade measurements taken per treatment reach. Some sites had five or six measurements, others had over twenty. Therefore the variability in the models should account for variability in the number of shade measurements. We performed a logit transformation of shade values to address to cope with the limited dependence variable (shade values cannot exceed 1.0) and as a means to linearize data that exhibited a curvature (see Section 1.6.2). Figure 3 illustrates the effects of such a transformation.

In the shade model-selection procedure, we considered models that included valley azimuth. These models did not perform well; the effects of valley azimuth on shade should have already been taken into account by our hemispherical photographs.

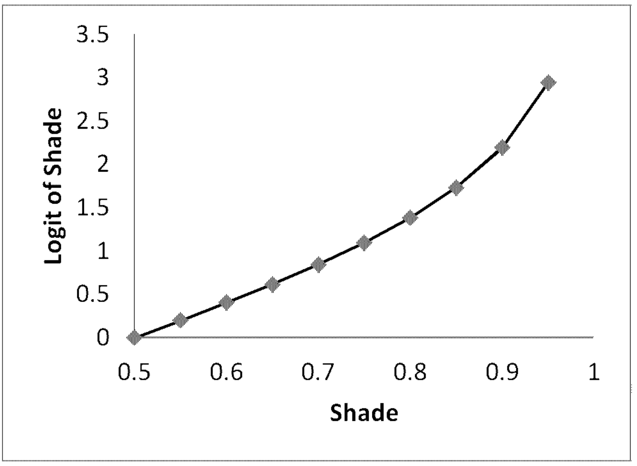

Figure 3. A display of the effects of performing a logit transform of the shade variable (shade values in this study fall between 0.5 and 0.95).

For logit-transformed shade data the best-supported statistical model was:

$$[2] \quad LogitShade_{Post} = \alpha_{Shade} + \beta_{1Shade}BasalArea_{Post} + \beta_{2Shade}TreeHeight_{Preharvest} + \beta_{3Shade}BasalArea_{Post} * TreeHeight_{Preharvest}$$

The logit of shade post-harvest depended upon the basal area post-harvest at each site, tree height values measured pre-harvest, and their interaction. The explanatory variables were assessed within 100 ft of the stream, a distance that we anticipated would include most if not all of the tree foliage that contributed to shading the reach during the middle of the day, when solar radiation is at its maximum. Basal area was calculated as the mean basal area (m^2/ha) of the two plots out to a horizontal distance of 100 ft, counting all living trees with diameters at breast height of 6" or greater. Mean pre-harvest height within 100 ft was 84 feet.

Main findings for the temperature modeling procedure included finding that shade (which, according to [1], changed values pre-harvest to post-harvest) is a major contributor in the model towards explaining model variation, with less shade associated with greater temperature increases between 2W and 3W (Table 1). The fixed value for Control was negatively associated with changes in the treatment reach temperature. The greater the change in stream temperatures between 1W and 2W, the greater an opposite change would be seen in the treatment reach. That is, if there were substantial cooling in the control reach in a given year relative to other years, the model would predict that even in the absence of harvest there would be a relative increase in the temperatures for the treatment reach. The power of the fixed and mixed-effects aspects of Control at influencing model fit is demonstrated in Figure 4. In particular, note that the cross-hairs (predicted temperature values) often are close to or essentially on top of the circles (observed values) even when the values are fairly far apart. This is an indication of the degree by which the model adjusts for year-to-year variation in stream temperature behavior.

Table 1 (From Table 4, Groom et al. 2011b). Fixed- and random-effect parameter values for a linear mixed-effects model and its associated temperature response variables. Parameters for treatment length and gradient are expressed as change in temperature per 1 km of distance or elevation. Observations = 119, Groups (Sites) = 33.

Fixed[a]	DF	Value	SE	p
Intercept	29.1	0.494	0.125	0.001
Control	21.5	-1.232	0.459	0.014
TRLength	28.2	0.800	0.304	0.014
Shade	94.5	-5.866	0.572	0.000
Grad1Q	30.3	-0.076	0.036	0.040

Random	Std.Dev
Intercept	0.441
ControlTemp	3.564
Residual	0.079

[a]Control reach temperature change = CT, gradient = GR, shade = SH, treatment length = TL.
[b]For Diel Fluctuation the variable for GR is replaced by elevation (EL). Other parameters in the model are the same.

Stream gradient (Grad1Q) and TRLength are associated with a change in temperature, with shallower gradients and longer treatment reaches associated with temperature increases over the course of the treatment reach. We interpret these findings to relate to the amount of time water within a stream is exposed to increased solar radiation in the harvest reach. Overall, when examining observed or partial residual temperature values for pre- and post-harvest periods, we found no temperature increase for State forest sites and ~ 0.7 °C increase for all private sites. A similar amount of temperature increase was predicted for a temperature model that used an indicator variable for ownership instead of a shade variable.

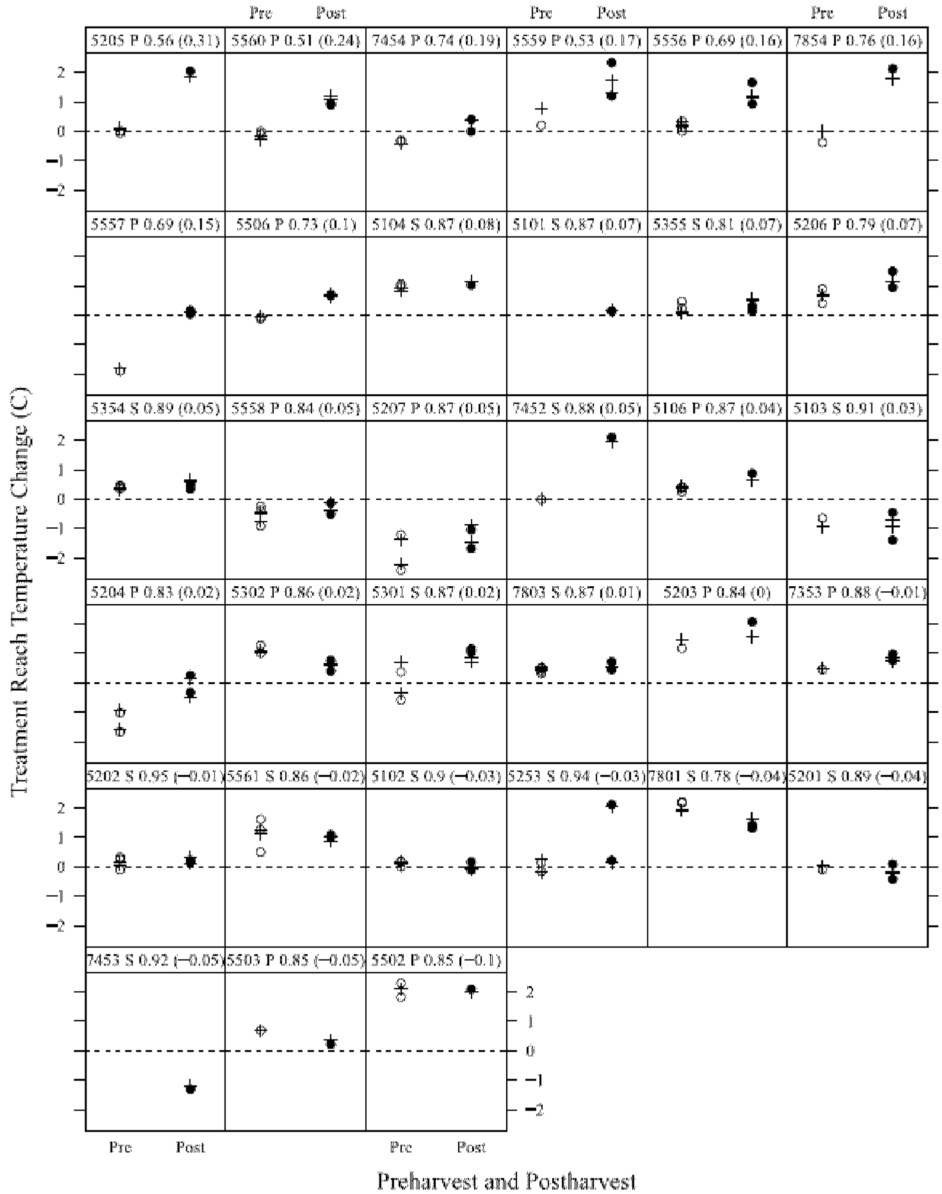

Figure 4, from Figure 3 in Groom et al. (2011b). Observed and predicted temperature changes for maximum temperatures (°C) by site. Pre-harvest and post-harvest observations are represented by open and filled circles, respectively. Each point represents one year of data collection at a site. The crosses represent predicted values from model Grad_Shade. Above each site's data is listed its site number, ownership ([S]tate or [P]rivate), post-harvest shade value, and in parentheses the change in shade value pre-harvest to post-harvest. Sites are ordered from the upper left to lower right by the observed change in shade values. Vertical differences of points within a pre-harvest or post-harvest category indicate a between-year change in the temperature relationship between 2W and 3W.

The shade model indicated that greater amounts of basal area post-harvest and shorter trees were related to greater amounts of shade. The model explained 0.69% (R^2 = 0.69) of the variation ($n = 33$, $\alpha_{Shade} = 1.795$, $\beta_{1Shade} = 3.100e^{-2}$, $\beta_{2Shade}= 6.250e^{-2}$, $\beta_{3Shade} = 4.680e^{-4}$, model $p < 0.001$). Other similar models were also explored.

1.5 Need for a more flexible approach – Bayesian modeling

The PCW analysis indicated that temperature increases above threshold levels were occurring on privately-owned timberland, and the effects analysis provided two models that linked temperature change within a reach to shade (eqn. 1), and shade values post-harvest to basal area and tree height (eqn. 2). We saw an opportunity to use the models from the effects analysis to predict the amount of basal area necessary to maintain shade levels, which in turn would be expected to prevent harvest-related stream temperature increases.

We needed a means to run the temperature and shade models backwards to reach predictions. That is, we wanted to provide basal area values for the shade model, obtain an expected shade amount, and then enter this shade amount into the temperature model to obtain a temperature change prediction. A major hurdle in performing such a procedure was that both of the shade and temperature model relationships contained error terms. We did not see a clear way to effectively join these models and rigorously account for error.

We became aware of a modeling approach called Bayesian Hierarchical Modeling. The idea behind this variety of modeling is that different relationships that affect each other can be combined, with findings from one relationship informing the other. An example of this approach involves Antarctic fur seals (*Arctocephalus gazella*; Hiruki-Raring et al. 2012). The authors determine factors associated with fur seal pup mass based on different relationships between prey densities, sea ice, maternal mass, and several other values. The relationships among these variables ultimately assist in estimating effects on pup mass. From our effects analysis, we have a temperature and shade model. The models do share a commonality: at least for post-harvest periods, shade appears as a dependant variable in the shade model and an independent variable in the temperature model.

Bayesian modeling in general is different from the modeling approach (Frequentist) used in the PCW and effects analysis. A central difference is the philosophical and analytical approach taken towards data analysis. In a Frequentist analysis data are assumed to be random expressions of fixed parameters (Ellison 1996). That is, if the world were virtually static and we could exactly repeat the same experiment many times, the processes affecting the data would be fixed, but like a coin flip the data would be similar, but not identical, to previous realizations of the experiment.

From a Bayesian perspective, it is the parameters that are treated as random and the data as fixed. The data resulted exactly the way they were because of myriad forces in the world. In the fixed-world example, if the experiment were run again the data would be identical. The various shifting effects in the world in turn make it so that our crude parameters of interest appear to follow random distributions.

There are some distinct advantages to using a Bayesian analysis.

1) Models may be hierarchically combined. That is, since our shade and temperature model share information (the Shade variable), the models could be joined to estimate values simultaneously while accounting for error associated with parameter fit and nuisance parameters.
2) Statistical models used in Frequentist analyses may be used for Bayesian analyses as well; it is the manner in which parameters are estimated that differ, not the structure of the parameters in the model. We therefore can make use of the shade and temperature models as formulated from the effects analysis.
3) Missing data may be estimated. Bayesian analysis works backwards and forwards simultaneously. While the analysis estimates dependent parameter values it can use the parameter estimates to create an estimate for the missing value.
4) Bayesian results more intuitive to understand than Frequentist results. Frequentist statistics speak to the probability of observed results given the repetition of data collection infinite times. A Frequentist 95% confidence interval tells us that the true value of interest (e.g., the true mean of a population) should fall within the estimated confidence interval 95% of the time if the study were repeated many times. A Bayesian 95% credibility interval (analogous to a confidence interval) around a mean indicates a 95% probability that the mean lies within the interval (Ellison 1996).
5) Bayesian models depend on incorporating knowledge. A Bayesian model must be provided with parameter priors; i.e., information regarding the size of an effect, its variance, and the expected distribution. If the magnitude of the effect is set to zero the prior is uninformative; Frequentist and Bayesian analysis results for the same model will be very similar if uninformed priors are used.
6) Obtaining derived estimates (e.g., predicted results given specified conditions) from a Bayesian analysis is relatively straightforward.

Bayesian analyses have some shortcomings:

1) Priors can influence model results if the priors are erroneous. This applies to the prior mean, variance, and distribution. If the variance is too restrictive or the distribution incorrect, the analysis may be prevented from correctly estimating values.
2) Convergence must be achieved. Bayesian analysis has only recently (10-20 years) become generally feasible as computer algorithms (i.e., Markov Chain Monte Carlo techniques) are relied upon to determine the posterior distribution of estimates run for $10^3 - 10^6$+ iterations. As the iterations run the parameter values improve (fit the data better). The results of the analysis are the various parameter point estimates taken from some number (e.g., 1000) of these iterations after the parameter values have stabilized. Each set of iterations for a parameter is called a chain. Modeling techniques allow us to assess the behavior of multiple chains. If one or more chains have not arrived at the same range of values then the model has not converged. We adjust the number of iterations to ensure that we include only values from chains that have reached convergence. At issue is that, particularly for complex models, the overall set of chains may have converged on a local value, not the "true" value.

3) If there are too many parameters in the model the model may overfit the data, or explain each datum well but have little ability to extrapolate well. Cross-validation techniques like leave-one-out can assist in assessing this condition.

We wished to use the Bayesian model to create model predictions. Specifically, we wanted to simulate harvests of different prescription types on our plot data and determine how the changes in the plot structure would affect shade values (via the shade model) and in turn how those changes would affect the change in temperature between 2W and 3W. We knew we could employ our best-performing temperature model from the effects analysis. However, for the Bayesian analysis, we decided to revisit the shade model, as the previously described model had some characteristics we wished to improve on (discussed in 1.6.2).

1.6 Developing sub-models

The temperature and shade models from the effects analysis represent different varieties of linear regression models. The temperature model included mixed effects parameters while the shade model is a weighted regression. We created null-prior Bayesian versions of each model type to ensure that resulting estimates matched the Frequentist estimates. Once we had the Bayesian parameterization decided upon, we could proceed with combining the models (Section 1.7).

1.6.1 Temperature

As a check of a correctly-functioning Bayesian model, we examined eqn 1 in R (library nlme, function lme, method = REML) as well as JAGS and WinBUGS. The R values differ slightly from Groom et al. (2011b), as the published values were run in SAS. The results of the probabilistic model are as follos:

```
              StdDev    Corr
(Intercept)   0.6683557 (Intr)
c_ControlTemp 1.8203856 0.306
Residual      0.2827235

Fixed effects: Response ~ Control + TRLength + Shade +
Grad1Q
                   Value Std.Error DF    t-value p-value
(Intercept)     0.476914 0.1246656 84   3.825543  0.0003
Control        -1.241839 0.4373531 84  -2.839443  0.0057
TRLength        0.828420 0.3006459 30   2.755467  0.0099
Shade          -5.923156 0.5647864 84 -10.487427  0.0000
Grad1Q         -0.083204 0.0423497 30  -1.964699  0.0588
```

Note that all independent variable values were centered for the analysis. The Bayesian results for the same model produced virtually identical findings for the fixed effects and similar estimates for the random effects (Table 2).

Table 2. Bayesian estimates for the effects analysis temperature model using uninformed priors. Percentages are available for constructing credibility intervals (95% CI = between 2.5% and 97.5%).

	mean	sd	2.50%	25%	50%	75%	97.50%
Random Effects							
intercept	0.71	0.11	0.52	0.63	0.7	0.77	0.96
Control	1.94	0.46	1.16	1.62	1.9	2.22	2.94
Residual	0.27	0.23	-0.2	0.12	0.28	0.45	0.68
	mean	sd	2.50%	25%	50%	75%	97.50%
Fixed Effects							
Intercept	0.48	0.13	0.21	0.39	0.48	0.57	0.75
Control	-1.24	0.48	-2.17	-1.55	-1.25	-0.93	-0.29
TRLength	0.83	0.32	0.23	0.61	0.82	1.04	1.51
Shade	-5.91	0.59	-7.07	-6.29	-5.9	-5.51	-4.74
Grad1Q	-0.08	0.05	-0.17	-0.11	-0.08	-0.05	0.02

The parameterization for the Bayesian model is provided in Appendix 2 (Section A2.1).

1.6.2 Shade

As mentioned in 1.5, we wished to re-visit the shade model. The effects analysis shade model was limited to examining post-harvest basal area out to 100 ft, not the full 170 ft measured. Therefore, our ideal shade model would:

1) Explain results well
2) Make intuitive sense
3) Include all vegetation plot data out to 170 ft horizontal distance from the stream
4) Include a measure of distance between stream and harvest at each site

The shade model from the effects analysis performed well and made sense (i.e., post-harvest shade was related to post-harvest basal area and tree height). However, it was limited to examining basal area only out to 100 ft horizontal distance. We originally limited it to 100 ft because the trees were on average 84 ft tall (less than 100 ft, greater than 75 ft); since solar radiation is most powerful between 10:00 AM and 2:00 PM in the summers, we anticipated that trees beyond this distance would not provide substantial shade to the stream. For this analysis we wished to more conclusively verify this assumption, and potentially include a measure of harvest distance to enhance the model's performance and inform the development of alternate harvest scenarios. Therefore we sought an analysis that incorporated all vegetation plot data recorded out to 170 ft and included a measure of harvest distance from the stream.

Incorporating a measure of distance from the stream to the harvest in the analysis proved difficult. We required a measurement that reflected actual harvest boundary distances. The vegetation plot data provide information on the distance of each tree from the stream as well as the vegetation plot line it was found in. To validate this method, we compared it against visually-determined harvest distances based on cumulative plots of basal area as a function of distance. By superimposing pre- and post-harvest cumulative basal area plot we observed where the two lines appeared to diverge. Selecting this divergence point was a subjective exercise, so we placed our assessments for every treatment plot in 25 ft categories (e.g., harvest between 26 and 50 ft horizontal distance; Figure 5). Assessments were corroborated by examining the difference between pre- and post-harvest values at 5 ft increments (not shown).

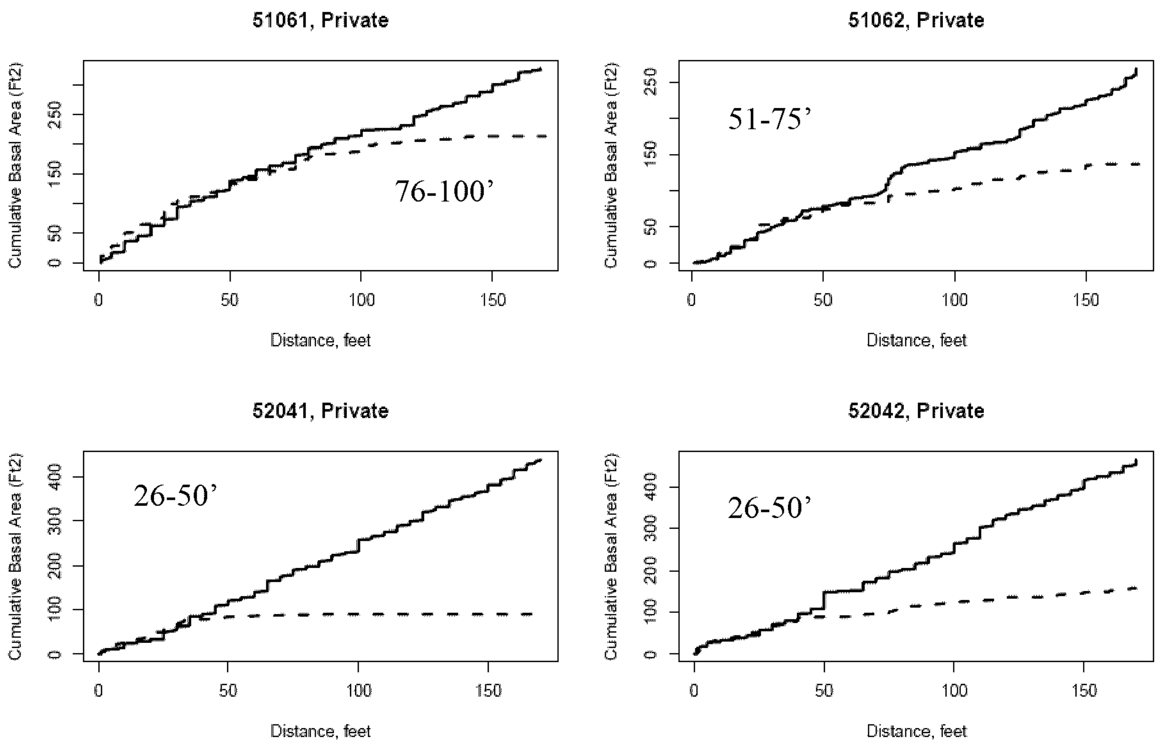

Figure 5. Cumulative basal area (ft^2) of four individual treatment reach vegetation plots, as a function of distance. Solid lines represent pre-harvest data, dashed lines are post-harvest. Observed departure points between the two lines are between 76-100 ft.

Once we had the visual assessments of harvest distances we compared them to the mean of the distance to the outermost trees in each of a vegetation plot's five lines (MeanMaxDist). We examined the relationship between the two harvest distance assessments to verify that MeanMaxDist approximated the visually-determined harvest distance (Figure 6). We interpreted the R^2 value of 0.88 as indicative of an essentially good fit. Therefore, MeanMaxDist could serve as a measure of harvest distance for our analysis.

We examined suites of models including basal area and MeanMaxDist. An interaction model including MeanMaxDist performed well:

$$[3]\ Shade_{post} = \alpha + \beta_1 BasalAreaPost_{170} + \beta_2 MeanMaxDist + \beta_3 BasalAreaPost_{170} * MeanMaxDist + \beta_4 TreeHeight_{Preharvest}170$$

Diagnostic plots indicated that it generally conformed to linear model assumptions, and had an R^2 of 0.71. The β_2 estimate was not significant but the interaction term (β_4) was. Although this model seemed promising, individual variables involved did not appear to exhibit a linear fit with shade (Figure 7).

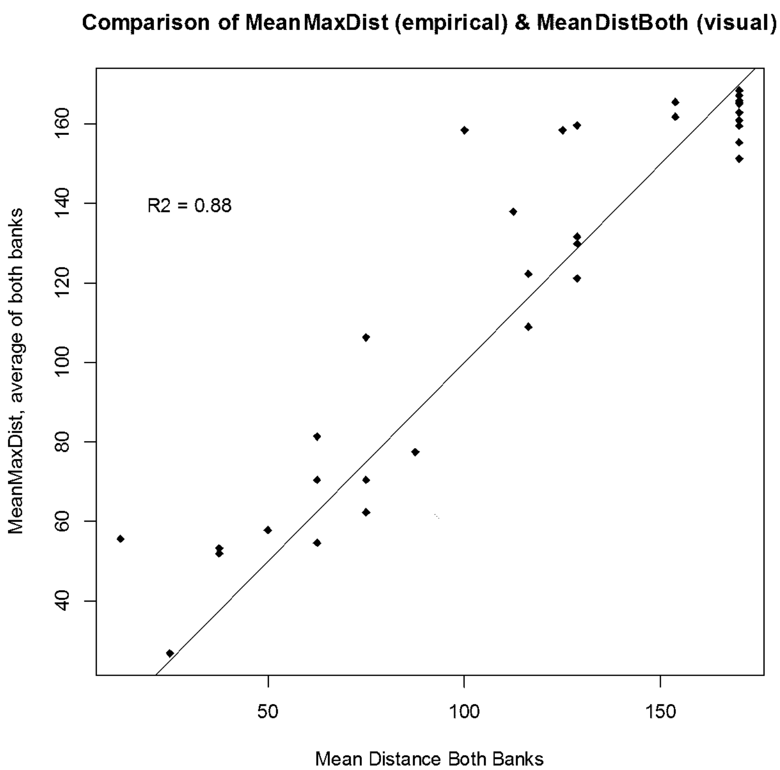

Figure 6. Comparison of vegetation plot distance empirical measurement (MeanMaxDist) and a visual assessment of harvest distance based on cumulative basal area by distance (MeanDistBoth). The line represents the linear fit of the two variables. Each point is the mean value of the two treatment vegetation plots at a site.

Shade appears to have a non-linear relationship with basal area of plots out to 170 ft and MeanMaxDist (Figure 7A, B). The relationship appears tighter for basal area and shade. This non-linear relationship encouraged a deeper consideration of how we were modeling shade. In particular, Figure 7B appears to have two different slopes on two different intervals in the relationship between post-harvest basal area and shade, with slopes changing at ~ 150 ft^2. Prior to that point the relationship is steep, afterwards less so. Our original assumption that trees may influence stream shading out to a certain distance from a stream appeared supported. We therefore examined a different way to consider the MeanMaxDist data in the analysis.

We suspected that harvests that were distant from the stream would have little impact on shade while those that were closer would have greater impact. At the same time, the closer the cut to the stream, the less basal area post-harvest. We therefore examined the relationship in Figure 7B by asking which of the basal area points corresponded to cut distances of </= 75 ft, 100 ft, and 125 ft from the stream (Figure 8). We indicate which points fell below or above the "cutpoints" and fit a line to each set of points. The slope of the green line (above the MeanMaxDist cutpoints) becomes shallower from Figure 8A to Figure 8C. The slope of the orange line (below the cutpoints) is virtually identical for Figure 8A and B and then becomes shallower and fits the below-cutpoints worse in Figure 8C. We interpret these figures to indicate that basal area information is most relevant for sites harvested within 100 ft of streams, but that the information quality is degraded if we look out to 125 ft.

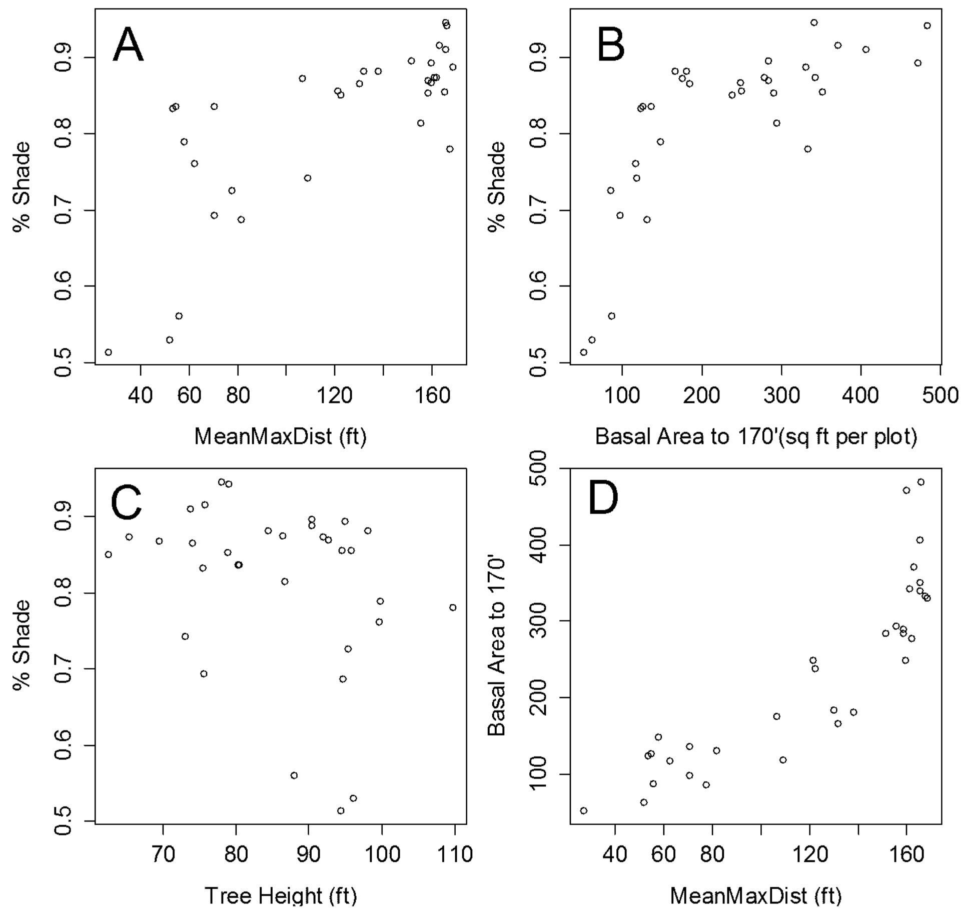

Figure 7. Scatterplots of percent shade post-harvest and MeanMaxDist (A), basal are of plots out to 170 ft(B), and pre-harvest tree height (C). The relationship between MeanMaxDist and basal area are plotted in (D).

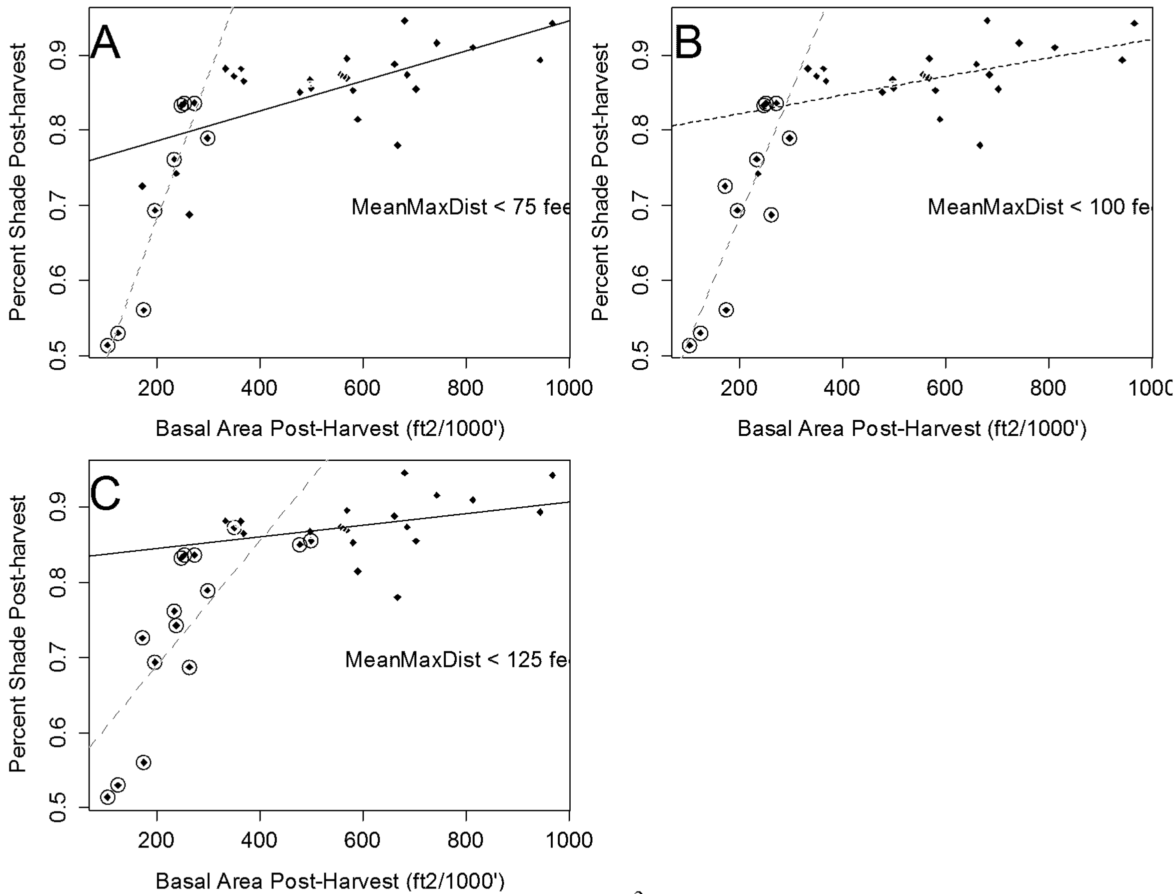

Figure 8. Fit of basal area post-harvest values (ft^2/1000 ft) to stream shade while parsing data by the distance to harvest from the stream. Figure 8A has circles around all site data where the variable MeanMaxDist </= 75 ft. Figures 8B and 8C do the same for MeanMaxDist </= 100 ft and 125 ft respectively. The solid black line is fit to points above the MeanMaxDist cutpoint, while the dashed grey line is fit to those data below the cutpoint.

We found further evidence that basal area beyond 100 ft was not useful for explaining shade. A line fit to the < 100 ft distance sites (Figure 8B) had a steep (and statistically significant) slope, indicating a relationship between shade and basal area post-harvest for sites with harvest distances < 100 ft. In contrast, the slope of the shallower (green) line is not statistically different from zero. A separate analysis that compared a model that limited inclusion of post-harvest basal area to within 100 ft of a stream to a model that included basal area to 100 ft and from 100 ft to 170 ft (two basal area measurements in the same model) indicated that the more simple model was preferable and that no information was gained by including trees from 100 to 170 ft (ΔAIC<2). We therefore interpreted these findings as justifying the inclusion of trees no further than 100 ft horizontal distance from the stream in the shade analysis.

With this decision we focused on shade model selection. Previously, in the effects analysis, we examined a model with the logit of shade as the dependent variable. Was this still supported? Comparing Figure 9 A and C to B and D (same independent variables, different dependent variables), it appears that there may be less of a curve evident in the data when shade is logit transformed and the independent variable is percent of basal area removed (Figure 9C). We therefore decided to use the logit of shade as the dependent variable. We used model selection to determine independent variables for inclusion. (Note: a quadratic term did not accommodate the curvature in 9B and 9D as well as the logit transformation.)

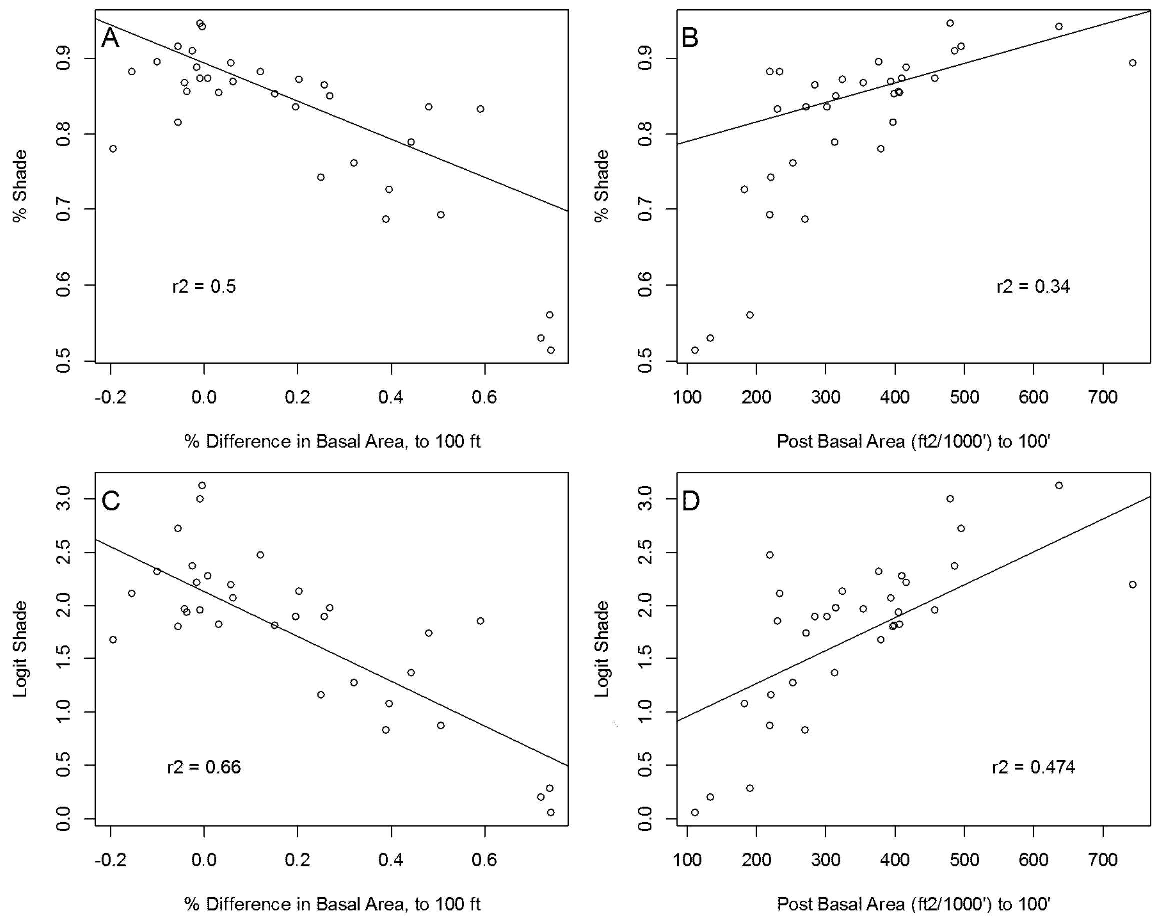

Figure 9. Plots comparing the fit of dependent variables shade (A, B) and the logit of shade (C, D) against independent variables "% difference in basal area, to 100 ft" (A,C) and "post-harvest basal area to 100 ft" (B,D). Lines represent weighted regression fits, R^2 values are provided.

We additionally wished to re-visit the metric for post-harvest basal area. Although post-harvest basal area originally performed relatively well as a predictor, it conveyed no information of pre-harvest stand information. Sites were generally well-shaded pre-harvest (figure 10). Yet, they differed in pre-harvest basal area values. We reasoned that a set reduction in basal area may disproportionately affect stream shading at sites with lower pre-harvest basal area values. Therefore we created the variable "Percent Basal Area Reduced," which is, for all trees within 100' horizontal distance, the amount of basal area pre-harvest minus the basal area post-harvest divided by the pre-harvest value.

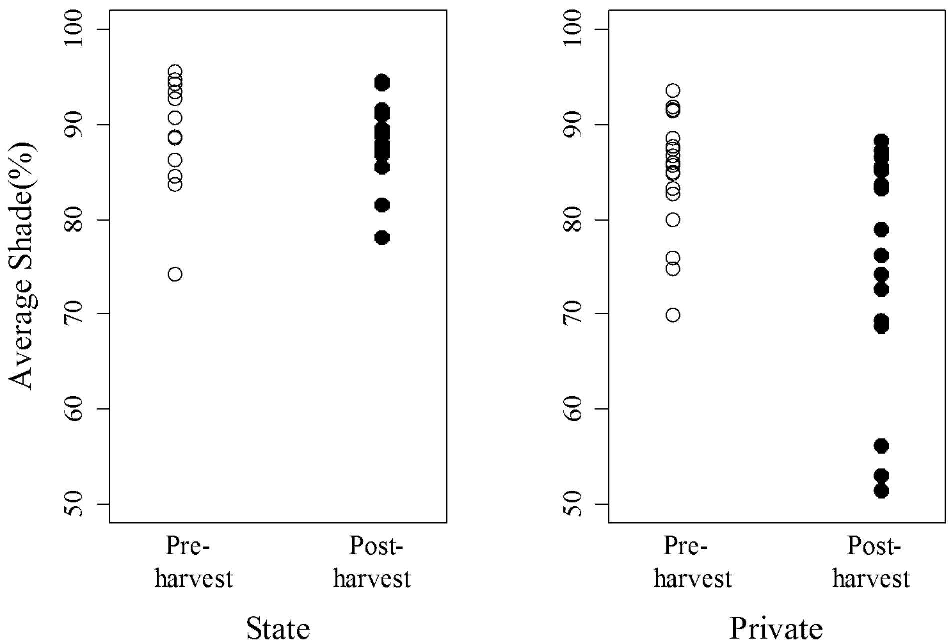

Figure 10. Appears as Figure 4 in Groom et al. 2011b. Plot of average treatment reach shade values (%) for each site, grouped by harvest status (pre-harvest, post-harvest). On the left are state forest shade values, on the right shade are private forest shade values.

For the independent variable model selection we investigated model fitting with combinations of the independent variables: post-harvest basal area to 100 feet, the percent basal area reduced, mean pre-harvest tree height of trees to 100 feet from the stream, and the percentage of hardwood pre-harvest within 100 feet, assessed by basal area. Models are presented in Table 3. The top six models included the variable "percent basal area reduced" which outperformed the variable for post-harvest basal area (present in the following six models). The top three models received similar AIC values (maximum ΔAIC = 2.53) and the probability that one of the three models were the best of the set was 0.99. Given that the third-best performing model had an AIC value slightly greater than the variable penalty term for the top-ranked model (which had one extra variable) we essentially see the explanatory power of all three premodels as roughly equivalent. Because we seek a generally robust model (fewest parameters) we select model 3 as the preferred model.

Table 3. Model selection for fitting logit of shade values. *k* is the number of estimable parameters in the model, *ΔAIC* is the change in AIC between a model and the model with the lowest AIC value, and ω is the model weight (probability of being the best of the set). Models are sorted by *ΔAIC*. All independent variables were assessed to a distance of 100 ft from the stream. PctBA_red is the percent basal area reduced, PostBA is the post-harvest basal area, Ht is average tree height pre-harvest, and PctHWD is the percent of hardwood by basal area pre-harvest.

Model	Independent Variables	*k*	*ΔAIC*	r^2	ω
1	PctBA_red + PctHWD + PctBA_red * PctHWD + Ht	5	0.00	0.81	0.56
2	PctBA_red + Ht + PctBA_red * Ht + PctHWD	5	1.46	0.80	0.27
3	PctBA_red + Ht + PctHWD	4	2.53	0.78	0.16
4	PctBA_red + PctBA_red^2	3	10.37	0.70	0.00
5	PctBA_red + Ht	3	11.17	0.70	0.00
6	PctBA_red	2	13.16	0.66	0.00
7	PostBA + Ht	3	13.34	0.68	0.00
8	PostBA + Ht + PctHWD	4	15.07	0.68	0.00
9	PostBA + PctHWD + PctBA_red * PctHWD + Ht	5	16.03	0.69	0.00
10	PostBA + Ht + PctBA_red * Ht + PctHWD	5	16.10	0.69	0.00
11	PostBA + PostBA^2	3	18.04	0.63	0.00
12	PostBA	2	27.41	0.47	0.00

The equation for the selected model for site *j*, where all variables are considered out to a distance of 100 ft from the stream, is:

$$[4]\ LogitShadePost_j = \alpha + \beta_1 PctBasalAreaReduced_j + \beta_2 PctHardwoodPre_j + \beta_3 TreeHeightPre_j$$

The observed vs. predicted fit for eqn 4 appeared essentially linear along the 1:1 line (Figure 11). Shade model diagnostics (Figure 12) indicate that residuals are fairly constant over modeled values, the data appear to be normally distributed, and that individual points are not exerting substantial leverage over the model. Points at the upper right of Figure 11 appear to be associated with greater standardized residuals relative to other points.

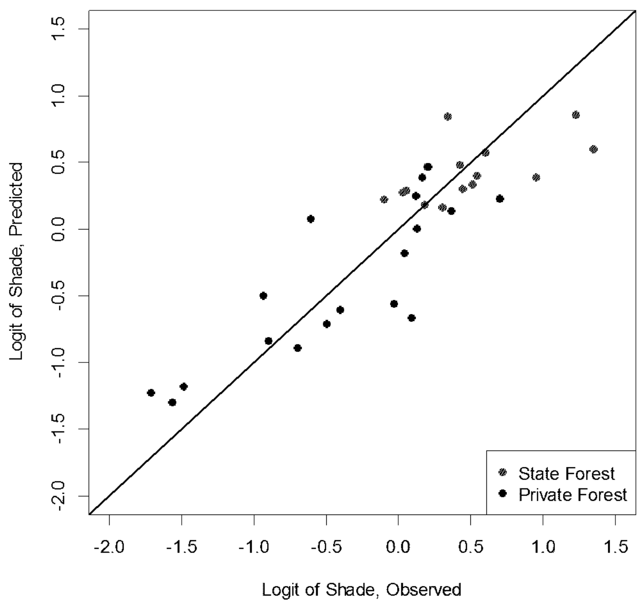

Figure 11. Weighted linear regression observed vs. predicted fit for the logit of shade (equation 4). The 1:1 line represents a perfect model fit; state forest sites are in red while privately-owned sites are in black.

The Frequentist fit for the estimates of the weighted shade model (Table 4) closely matches the Bayesian fit (Table 5). The measures of error differ, which appears to be an effect of the inclusion of sample weights. The parameterization for the Bayesian shade model is provided in Appendix 2 (Section A2.2).

Table 4. Output from R for modeling the logit of shade according to eqn. 4. The regression is weighted, with weights = 1/(variance in logit shade).

	Estimate	Std. Error	t value	Pr(>|t|)
Intercept	-0.02507	0.05111	-0.491	0.62742
PctDiffBA	-2.30961	0.23724	-9.735	1.21e-10
TreeHt	-0.04406	0.01515	-2.909	0.00690
PctHWD	-0.74600	0.21445	-3.479	0.00161

Residual standard error: 0.5077 on 29 degrees of freedom

Table 5. Bayesian estimates for the shade model described in eqn. 4 using uninformed priors. Percentages are available for constructing credibility intervals (95% CI = between 2.5% and 97.5%).

	mean	sd	2.50%	25%	50%	75%	97.50%
Intercept	-0.02532	0.10027	-0.22416	-0.09287	-0.0254	0.044207	0.166609
PctDiffBA	-2.31089	0.468524	-3.22638	-2.62181	-2.31706	-1.99762	-1.38907
TreeHt	-0.04382	0.029668	-0.10029	-0.06417	-0.04407	-0.02384	0.014772
PctHWD	-0.73969	0.421963	-1.55758	-1.0295	-0.74402	-0.4573	0.108633

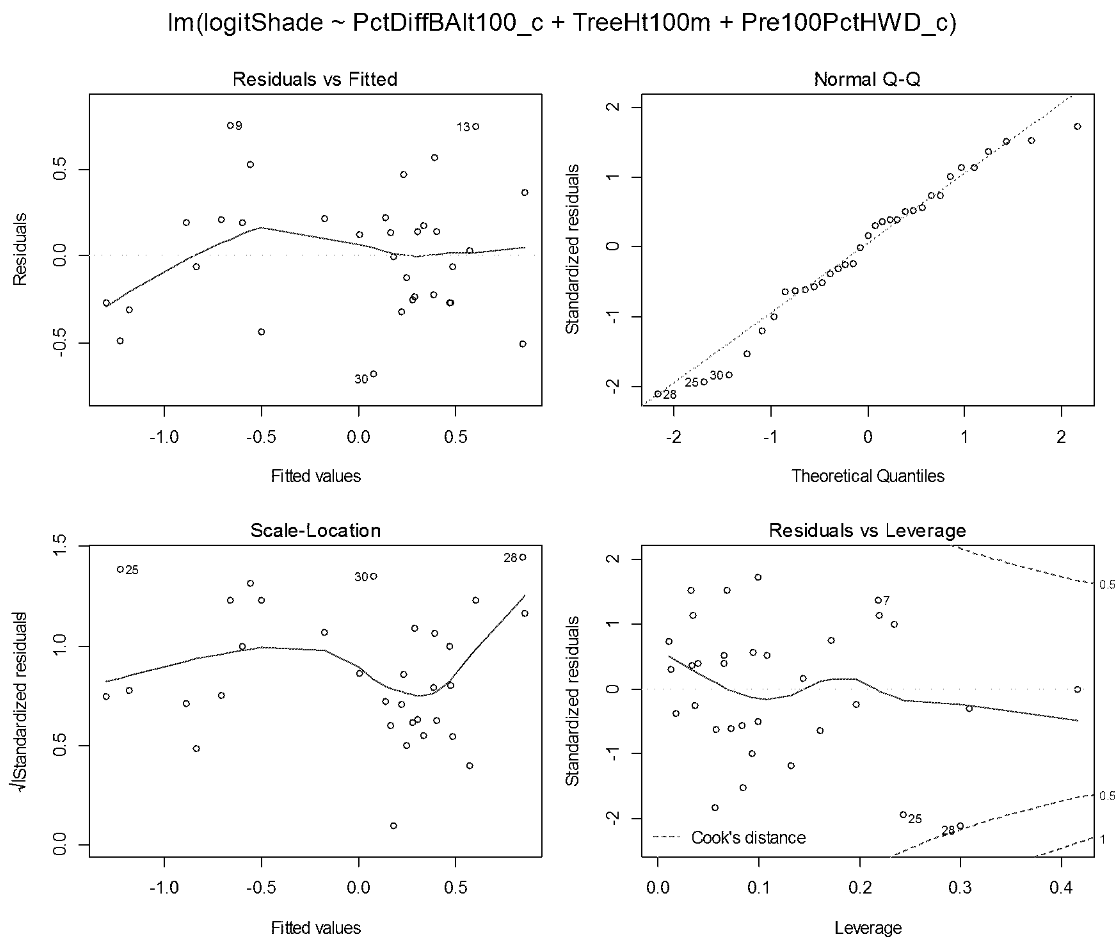

Figure 13. Diagnostics plots for the selected logit shade regression model (eqn 4).

1.7 Combining the sub-models

Linking the shade and temperature models was relatively straightforward. Our goal was to estimate post-harvest responses of shade and stream temperature to specified harvest prescriptions. Therefore, we did not need to alter the temperature model or its inputs for the pre-harvest period. Observed shade values were used as shade variable values pre-harvest. For the post-harvest period shade variable values were replaced with the *estimated* post-harvest shade values (eqn. 4). This replacement was done to link the temperature sub-model and the post-harvest shade sub-model. The post-harvest shade sub-model itself was informed by post-harvest shade data. Both the observed and estimated logit shade values were transformed to percent shade (Appendix 2, Section A2.3).

[5]

$$\Delta T_{ij,k=1} = \alpha_0 + \alpha_j + (\beta_1 \Delta TControl_{2-1ij} + \beta_j \Delta TControl_{2-1ij}) + \beta_2 TRLength_j + \beta_3 Shade_{j,k=1} + \beta_4 Grad1Q_j$$

$$\Delta T_{ijk=2} = \alpha_0 + \alpha_j + (\beta_1 \Delta TControl_{2-1ij} + \beta_j \Delta TControl_{2-1ij}) + \beta_2 TRLength_j + \beta_3 Shade_{Post} + \beta_4 Grad1Q_j$$

$$LogitShadePost_j = \alpha_{Shade} + \beta_{1Shade} PctBasalAreaReduced_j + \beta_{2Shade} PctHardwoodPre_j + \beta_{3Shade} TreeHeightPre_j$$

In Eqn. 5 ΔT is the change in stream temperature between 2W and 3W (3W – 2W); i = year, j = site, and k = timing relative to harvest (pre- or post-harvest). For clarification, where $k = 1$, $i = 1, 2, n$ years pre-harvest, where $k = 2$, $i = 1, 2$ years post-harvest.

Eqn. 5 contains all of the non-nuisance parameters estimated by the combined Bayesian model. Of note, parameter values in the shade sub-model have a *Shade* suffix, and we perform a logistic transformation of logit shade values and estimates for use in the temperature estimation equations. The temperature change sub-model maintains the same parameters pre- and post-harvest ($k = 1, 2$); therefore, their parameter distributions are estimated using data from both the pre- and post-harvest periods. Shade sub-model estimates are informed by post-harvest shade data and the temperature sub-model estimates by observed changes in temperature. This involves estimations of model error terms and other nuisance parameters (Appendix 2, Section A2.3). We use these estimated parameter distributions for the prediction scenarios.

Non-nuisance parameter estimates for [5] are presented in Table 6. Parameter estimates for the full model are similar but not identical to those presented in Tables 2 and 5. Of note, the credibility intervals for shade values have become narrower. For a more complete listing of parameter estimates, see Appendix 2, Section A2.4.

Table 6. Parameter estimates (abbreviated) for estimates of [5]. Sub-models are underlined. Sub-model means, standard deviations, and quantiles are presented.

Sub-models & Parameters	mean	sd	2.50%	25%	50%	75%	97.50%
<u>Temperature</u>							
Random Effects							
Control	1.921	0.440	1.186	1.599	1.887	2.199	2.861
Intercept	0.730	0.111	0.538	0.654	0.720	0.794	0.981
Residual	0.153	0.228	-0.297	-0.007	0.165	0.319	0.555
Fixed Effects							
Intercept	0.396	0.133	0.104	0.310	0.400	0.484	0.645
Control	-1.092	0.460	-1.951	-1.399	-1.105	-0.807	-0.152
TRLength	0.871	0.336	0.206	0.663	0.878	1.077	1.511
Shade	-5.606	0.844	-7.341	-6.153	-5.590	-5.030	-4.046
Grad1Q	-0.077	0.049	-0.179	-0.109	-0.076	-0.044	0.014
<u>Shade</u>							
Intercept	-0.279	0.066	-0.407	-0.321	-0.279	-0.237	-0.148
PctDiffBA	-2.776	0.305	-3.428	-2.973	-2.770	-2.558	-2.223
PctHwd	-0.585	0.249	-1.092	-0.754	-0.583	-0.414	-0.100
TreeHt	-0.065	0.017	-0.100	-0.076	-0.066	-0.054	-0.031

We are interested in predicting the effects of specific harvest scenarios on the changes in stream temperature for treatment reaches. We do so by controlling as many variables as possible. Our approach is to predict a change in stream temperature using a harvest scenario resulting in a specific per-site value for Percent Basal Area Reduced (PctBasalAreaReduced).

$$[6]\ \Delta\hat{T}_{i=1,j,k=2} = \alpha_0 + \alpha_j + (\beta_1 \Delta TControl_{2-1ij} + \beta_j \Delta TControl_{2-1ij}) + \beta_2 TreatmentReachLength_j + \beta_3 (inverse\ logit\ of{:}\ \alpha_{Shade} + \beta_{1Shade} \boldsymbol{PctBasalAreaReduced_j} + \beta_{2Shade} PctHardwoodPre_j + \beta_{3Shade} TreeHeightPre_j) + \beta_4 GradientQuartile_j$$

Equation [6] is populated with the estimated parameters from [5]. Importantly, change in temperature for equation [6] represents a derived value, not an estimated value. The Bayesian model obtains all estimated values from [5]. For every scenario, we obtain two predicted temperatures from [6]. The first is the predicted change in treatment reach temperature for the first year post-harvest with a harvest effect (simulated or observed). The harvest effect is represented by the variable *PctBasalAreaReduced* and is calculated from vegetation plot data (see Section 1.2.3). The second prediction sets the variable *PctBasalAreaReduced* equal to zero change between pre- and post-harvest basal area. We subtract the second prediction from the first, with the difference representing the predicted increase in temperature, per site, due to harvest. This procedure effectively controls for site-specific influence of non-shade temperature variables.

For a formulation of the model and incorporation of temperature prediction, see Appendix 2, Section A2.3.

1.8 Model evaluation

There were several aspects of model evaluation that we investigated. We wanted to determine model fit (i.e., how well did it predict observed values) and determine how well the model met assumptions. These evaluations do not consider simulated results because there are no observed values or "truth" against which to directly check predictions.

For model fit, we examined the first-year post-harvest estimated values as these were the values we wished to simulate. We plotted estimated vs. observed values (Figure 13). Of note, the estimated values represent the change in temperature between the probes 2W and 3W, taking into account all variables in the model including Percent Difference in Basal Area. We interpret the plot as indicating a linear fit with some under-estimation of temperature change at the larger increases in temperature.

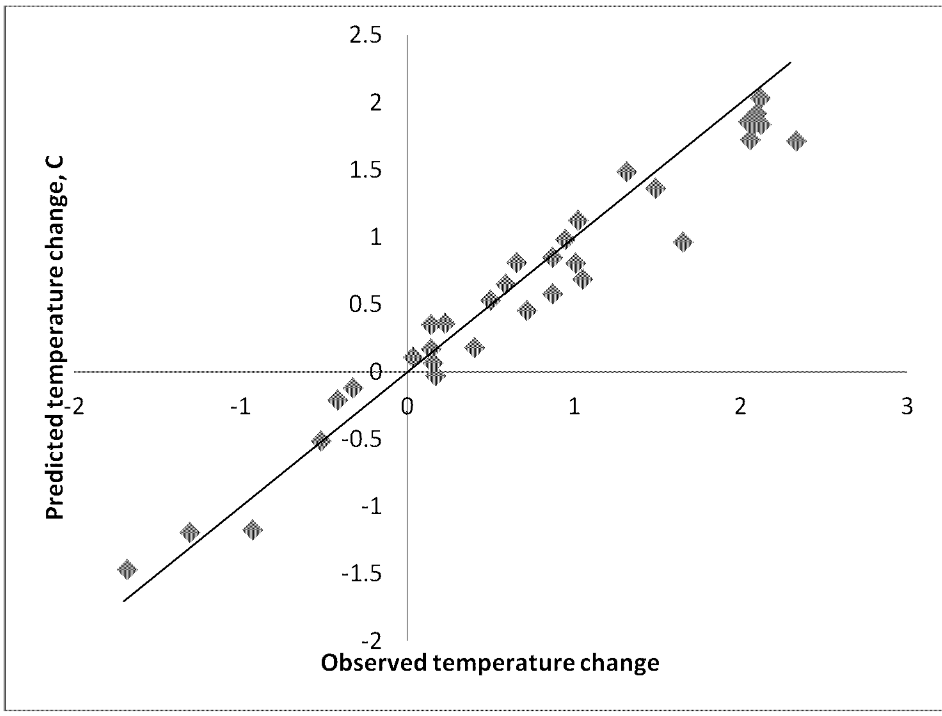

Figure 13. Observed vs. predicted changes in treatment reach temperatures between probes 2W and 3W. The diagonal line represents a one-to-one fit; the blue diamonds represent the change in temperature between 2 and 3 W for the first year post-harvest.

We then examined model assumptions, some related to predicted values.

1.8.1 Assumption: Priors specification

We performed two evaluations of priors. The first was to examine the distributions of priors and determine if any appeared to have distributions truncated by prior specification. All prior distributions appeared to function well except for the uniform distribution for the nuisance parameter idelS (Appendix 2, section A2.2). As a uniform distribution, it estimated a value precisely at the distribution's midpoint (50 if from 0 to 100). We increased the possible range of values to 0-1000, and found that it predicted a mean of 500. Because this parameter was for the weighted shade regression model, we examined replacing the uniform distribution with a gamma distribution and checked model performance. We executed the shade analysis with the gamma distribution and null prior values for the mean and variance of 0.01 and 0.001. All four of these iterations produced very similar outcomes in other parameters (which in turn were approximately equivalent to the Frequentist shade model). We interpreted the original prior as functioning well.

The model contains many priors with Gaussian distributions. They all have starting means of 0, but typically have precisions (reciprocal of variances) of 0.001. We increased the variance to 0.01 for all parameters and ran the model for 2×10^6 iterations. Results were very similar between the two levels of precision. Most means of parameter distributions exhibited less than a 1% difference; the largest absolute difference was for the covariance parameter estimate (4%). We interpreted these results to indicate that the priors were appropriately specified.

1.8.2 Assumption: Parameter convergence is achieved

Models were run with 6 chains (i.e., 6 parallel estimation procedures) for a number of iterations beyond which we saw any steps in trace plots for any of the parameters. We cannot know if parameters settled for all six chains reliably at the same range of values but would have settled at different values for longer and longer runs. We were careful to include sufficient iterations to reach an estimation consensus.

1.8.3 Assumption: Model is not overfit

An overfit model is one which lacks generalization; it describes the data at hand well but would fail to predict useful responses given new inputs. One method for assessing this condition is to conduct a leave-one-out cross-validation. The idea is to see how well the model predicts a data point if that point were not included in the analysis. We conducted this analysis to re-create first-year post-harvest estimates in a similar fashion to Figure 13. Individual first-year post-harvest values were omitted. The Bayesian model imputed the missing values using the within- and among-site estimates and relationships.

An important consideration for the cross-validation approach is that we are conducting it for a mixed-effects analysis. When we leave out a data point, it affects the estimation of the random effects parameters for a particular site. This is especially true for sites where there may be a single value pre-harvest and post-harvest due to probe malfunctions. In these instances the mixed-effects parameter may resort to a global mean, altering the prediction greatly.

The leave-one-out cross-validation results are presented in Figure 14. The points are not as linear at higher temperature changes as in Figure 13. To an extent this should be expected, with less information available (particularly for random effects that are using at the most four or five data points to begin with). We have identified the data points from sites that were reduced to one point per site in the leave-one-out (LOO) procedure (red circles). The three biggest outliers are for an observed decline of 1.3 ^{0}C and observed increases of 2.1 ^{0}C. We interpret the LOO to indicate that the model sufficiently fits the data. Extreme outliers were generally those that had only one data point to use for random effects estimations, and if anything the model indicates an underprediction of temperature increase (conservative) relative to observed temperature increases.

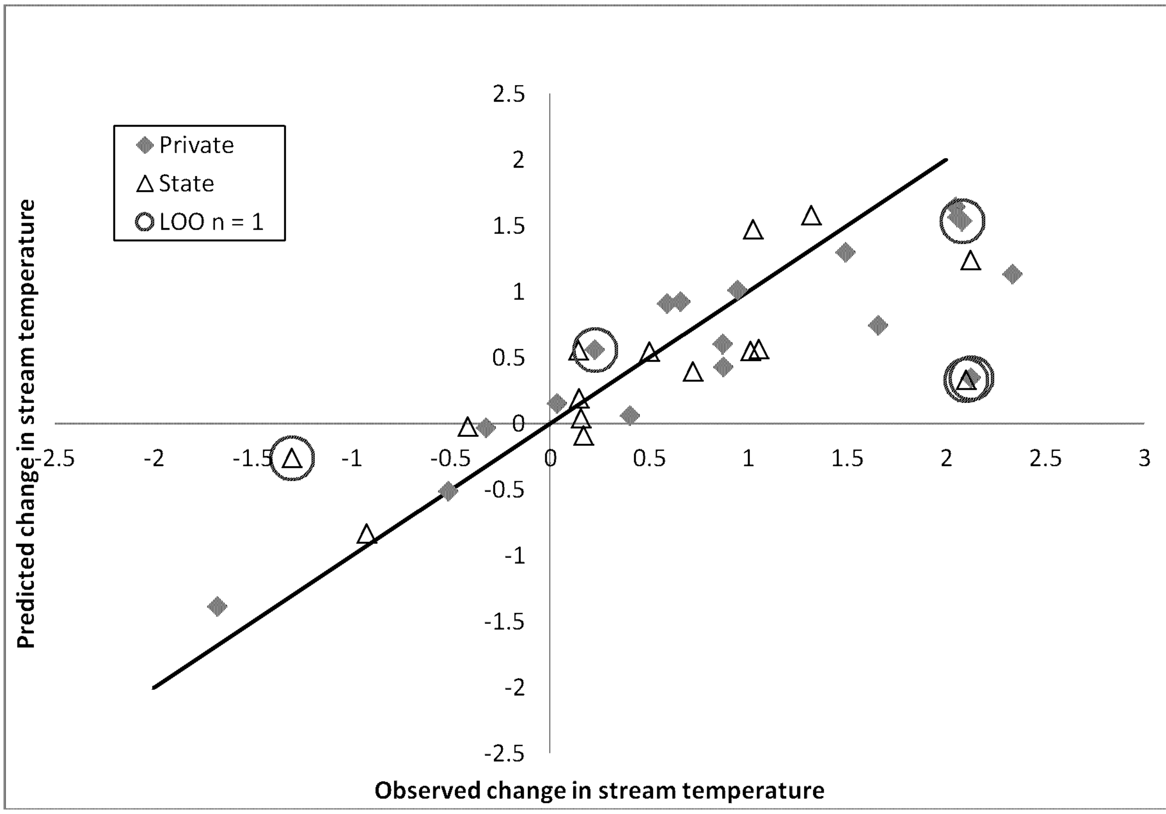

Figure 14. Leave-one-out cross-validation of Bayesian model, first-year post-harvest data. Each point represents an estimated temperature response given the remainder of the temperature data set. Open triangles are state forest sites; closed diamonds are privately owned sites. Red circles surround data points that were estimated with only one data point informing random effects.

1.9 Model – specific assumptions

We performed assumption-checking procedures in Section 1.8 to verify the validity and performance of the model. Those procedures were standard for evaluating a predictive Bayesian model. However, there are some important assumptions for this specific model that affect interpretation and use of the model.

1.9.1 Inference and utility

Our model appears to fit the data well (Figure 13). However, its data requirements make it unlikely to suitably apply to other data sets. The model utilizes pre-harvest and post-harvest data for a treatment and control reach (with the control reach directly upstream of the treatment reach). It also depends on shade data and riparian vegetation information. It requires these data to feed a specific mixed-effects temperature sub-model and the associated shade sub-model. Therefore, we do not foresee the model being used directly outside of these study sites.

In addition, sites were not randomly selected. We used virtually all sites provided by industry and state forests that met our site inclusion criteria (including agreement to maintain an unharvested upstream control reach). Given our sample size (33 sites) and geographic extent of sites we interpret the results as representative for small and medium type-F streams in the Oregon Coast Range (three sites technically fell in the Interior georegion). Technically we lack the statistical inference a randomized study offers. On the other hand we are not aware of other randomized stream temperature studies that offer the same level of statistical power as this manipulative long-term study design.

1.9.2 Pre-harvest shade levels

Our model describes shade and temperature relationships for sites that all began with levels of shade that typically exceeded 80% (Figure 10; 4 sites had shade levels pre-harvest between 70 and 80%). It therefore may not describe well the thermal behavior of streams at sites subject to harvest that exhibit lower pre-harvest shade levels.

1.9.3 Harvest tree distributions

As described in Section 1.6.2 we excluded consideration of trees (basal area) measured farther than 100 ft from the stream, horizontal distance. However, within 100 ft of the stream we relied on percent basal area reduced, not the mean of the maximum distance to the edge of the harvest. Twenty-two sites had values for the mean of the maximum distance to the edge of harvest beyond 100 ft, so the variable was relatively non-informative when compared to percent difference in basal area.

Only one state forest site had substantial reduction in basal area within 100 ft of the stream; therefore, reductions in basal area predominantly reflected private ownership harvests. These harvests appeared as hard-edged clearcuts (i.e, not thinnings within the RMA). The variable for the percent difference in basal area within 100 ft of streams is therefore reflective of a hard-edged clearcut and not thinning. It is informed purely by changes in basal area within 100 ft from the stream; therefore, we assume in our simulations that harvest occurs from 100 ft from the stream and inwards which is reflective of our observations of site harvest patterns.

Part II: Harvest Simulation Approach

2.1 Vegetation data use

Vegetation plot data collection is described in Section 1.2.3. To summarize, we have 100% tree cruise information from all plots pre-harvest, from treated plots post-1, and blowdown were tallied in all plots during the post-1 period. We summarized the vegetation data from each plot to obtain different metrics (MeanMaxDist, basal area pre-harvest and post-harvest) for use in the shade analysis and to develop the metric Percent Basal Area Reduced.

To simulate a harvest we used pre-harvest plot data. These data were then subject to a specified harvest procedure (e.g., FPA harvest). We ran identical plot summary programs on the pre-harvest plot and the simulated harvest plot. One of the variables recorded for both scenarios (pre and post) was basal area within 100' horizontal distance of the stream. The horizontal distance is used in all cases as it is the metric the predictive analysis relies upon.

Plot vegetation data are further summarized into site data. Site vegetation characteristics are the mean of the corresponding treatment plot metrics. Therefore, if a simulated FPA harvest removed more trees from one bank than the other, the resulting value for basal area would fall between the two plot basal area values. We performed this step as the model itself was constructed with the vegetation data summarized in this fashion. As mentioned in Section 1.4, the model does not incorporate valley azimuth.

To obtain Percent Basal Area Reduced we subtract a site's post-harvest basal area (the mean of basal area values from the two treatment plots) from the site's pre-harvest basal area, divided by the pre-harvest basal area. This procedure was used to predict the effects of harvest under all scenarios except the as-harvested scenario. In the as-harvested scenario we used both pre-harvest and the measured (not simulated) post-harvest data.

Obtaining the measured post-harvest data for the as-harvested scenario required additional data manipulation. As mentioned above, only treated plots were fully re-measured post-harvest. If a site was harvested only on one side, we combined the treated plot with the pre-harvest data from the untreated plot. A further complication was blowdown. For estimating the effect on stream temperature, we removed plot blowdown from the basal area estimate. We omitted blowdown as it would likely reduce shade levels and we wished to predict as closely as possible the temperature increase due to the change in shade.

2.2 Predictions

2.2.1 As-harvested

As described in Section 1.7, we used the observed change in basal area between pre- and post-harvest to predict harvest-related changes in temperatures for individual sites (Figure 15). The privately-owned sites had predicted overall temperature increases of 0.93 ^{o}C while the state forest sites had predicted increases of -0.05 ^{o}C.

Within this prediction some of the State sites exhibited predicted temperature decreases. This was due to greater basal area within 100 ft of streams recorded post-harvest compared to pre-harvest vegetation cruise surveys. We attribute the increase in basal area due to tree growth, ingrowth of smaller trees, and measurement error.

2.2.2 State Forests

We developed an approach to simulate a state forest FMP riparian harvest of our pre-harvest treatment reach data (Appendix 3). All treatment reach pre-harvest vegetation data were reduced according to our interpretation of the FMP, simulating harvest on both banks (Figure 16). Sites that achieved mature forest condition, either by being dominated by hardwoods or by having many large conifers, were not reduced in basal area by the simulation (temperature increase = zero). Other sites received harvest within 100 ft as limited by minimum basal area and conifer number retention requirements. All other possible trees were removed. The analysis predicts that the average temperature increase for all sites subjected to a thorough FMP harvest is 0.19 ^{o}C, with a 95% probability (one-sided) that the mean would be less than 0.23 ^{o}C.

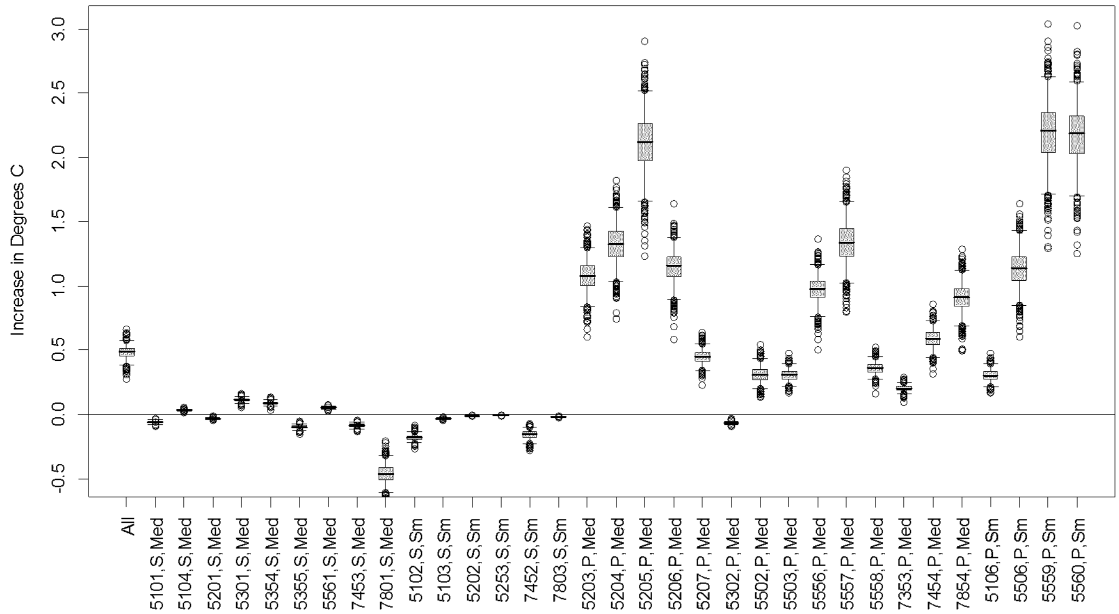

Figure 15. Predicted temperature increase due to change in basal area for all sites, as harvested. Boxplots indicate median values (mid-lines) within grey boxes (25% and 75% quartiles), with whiskers extending from the $2.5^{th} - 97.5^{th}$ percentile. Circles represent data outside of the whiskers. Sites are listed along with ownership (S or P for State or Private) and stream classification (Sm = small, Med = medium). The horizontal line indicates zero temperature increase.

2.2.3 FPA Harvest

We programmed a similar two-sided harvest of all sites, irrespective of ownership, to experience a complete removal of as many trees as permitted by the FPA (Appendix 4). This harvest relied on tree slope distance from streams. We then determined the percent change in basal area between the pre-harvest plots and the simulated FPA harvests of those plots and entered the value into the simulation portion of the Bayesian model. The results indicate a mean temperature increase of 1.77 °C with only a 5% chance that the temperature increase on average would be below 1.43 °C (Figure 17). This result indicates that 1) removal of all trees permitted under the FPA is predicted to result in significant warming, and 2) given the discrepancy between the as-harvested predicted warming and the simulated FPA harvest, a harvest as extreme as is portrayed in Figure 17 may not be common practice on industrial land ownership. We initially asked landowners participating in the RipStream study to extract timber from riparian areas as permitted by the FPA and NWFMP. This was done to ensure that the rules, not practices, were tested. We were concerned that a "business as usual" harvest would not allow us to test the rules (according to ODF's 2002 compliance audit[1] 60% of landowners did not enter the RMA and on average operators left over 200% of the required basal area). Even with landowner

[1] ODF 2002. Best Management Practices Compliance Monitoring Project: Final Report. Oregon department of forestry forest practices monitoring program technical report #15.

cooperation, it appears that harvestable basal area may have remained.

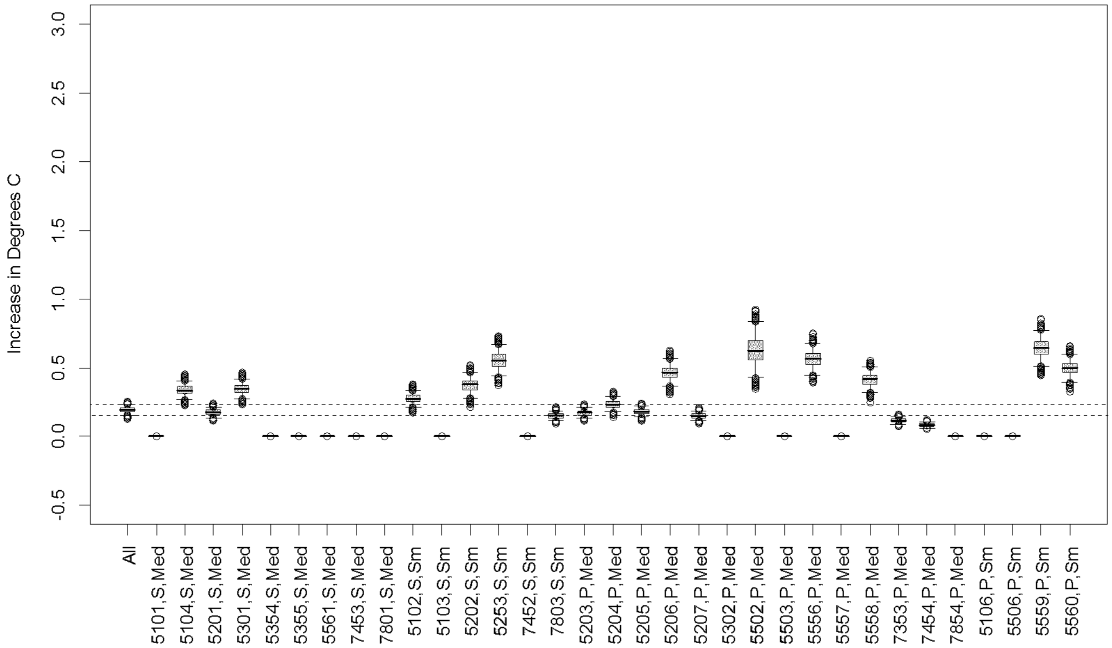

Figure 16. Simulation of response of all sites to a two-sided State Forest FMP harvest. The "whiskers" (lines above and below boxes) identify the 95% credibility interval for each site (95% probability that the mean response falls between those lines). The boxplot on the far left, labeled "All," represents the mean response of all 33 sites, and its 95% credibility interval is represented by the two dashed lines.

Riparian prescriptions in the FPA are to abide by rule language *averaged over 1000 feet.* Therefore, it is possible that the riparian plots, which are each 500' long, captured at some sites metrics not representative of buffer dimensions or characteristics. Ancillary work examining orthophoto imagery and LiDAR on suspected aberrant sites provided no indication that vegetation plots were unrepresentative. Figure 18 displays the total basal area from pre-harvest data, as-harvested data, and the FPA and NWFMP simulated harvests. It appears that state forest and privately-owned sites typically harvested less than they potentially could have. All but one state forest site harvested well above our simulated NWFMP level. Approximately six out of the 18 private sites appeared to harvest at or almost at the FPA level. Twelve of the 15 State Forest sites appeared to essentially receive little or no entry within 170' of the streams at the location of the vegetation plots.

The RipStream vegetation plots were cruised by forestry professionals who were experienced at conducting state forest Stand Level Inventory cruises. Additionally, the Private Forests monitoring team performed quality checks of every site each time a vegetation plot cruise was completed. We therefore believe these data are credible.

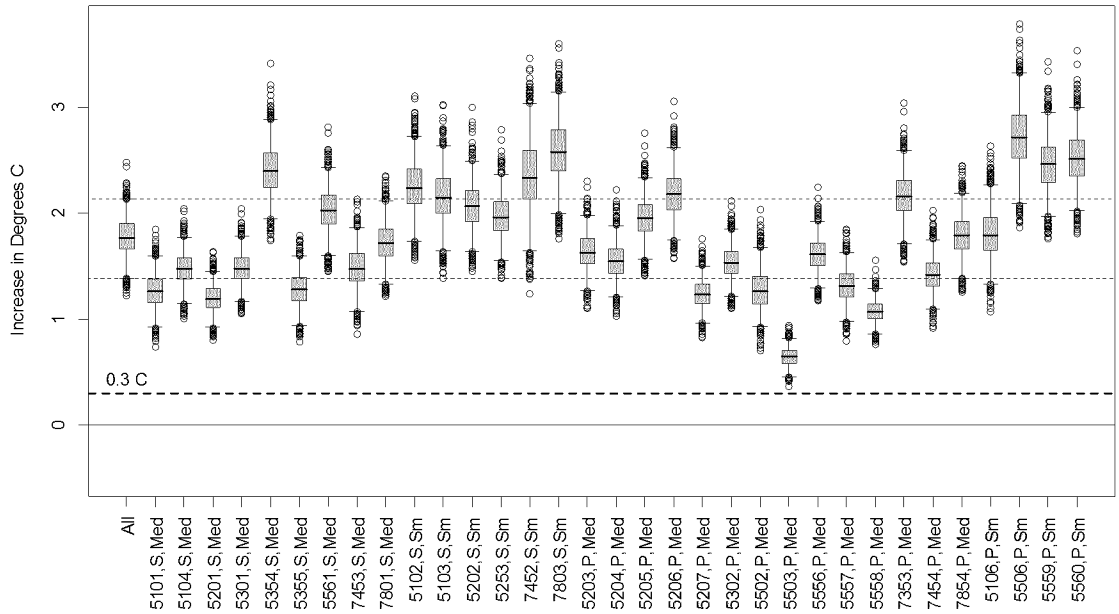

Figure 17. Predicted temperature increases from simulated FPA harvest of all sites. The leftmost boxplot, "All," represents the mean response of all sites. The two dashed lines above 1 °C represent the 95% credibility interval for All, indicating the 95% probability that the mean response would be within that band. The other dashed line indicates 0.3 °C.

2.2.4 Percent Harvest

Prior to developing and testing prescriptions we wished to test the model's performance against a suite of conditions. Since the model relies on changes in the variable *PctBasalAreaReduced* we performed an examination of anticipated temperature increases with incremental changes in this variable. The predicted temperature increase changed as each site's basal area was reduced in 10% increments from 80% to 10% of pre-harvest values (Figure 19). Although the sites exhibit different temperature responses to basal area removal (Figure 19A) it appears that, when examining basal area removal within 100 ft horizontal distance of the stream, the 0.3 °C threshold is generally crossed at around a 15% removal of basal area (orange line in 19A, mean line in 19B). The mean lines indicate a 50% probability that the mean response will be below or above that line at the given level of basal area reduction. The credibility interval for the mean (Figure 19B) indicates that the probability of the mean response crossing 0.3 °C lies between a 12 and 18% basal area reduction.

Harvest Comparison

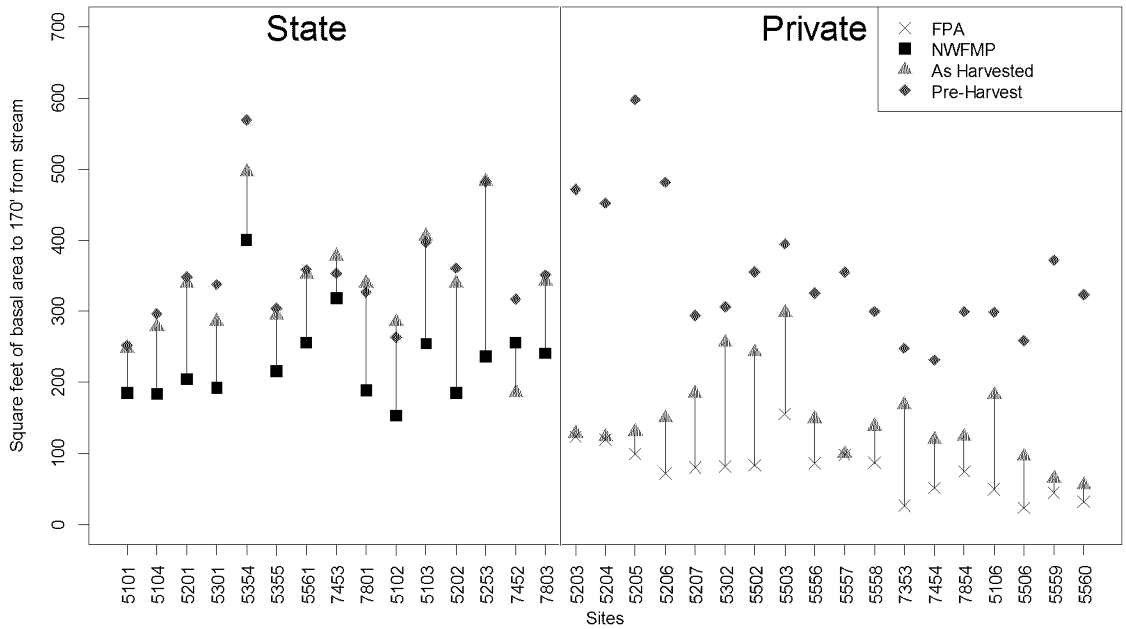

Figure 18. Square feet of basal area from 170 ft x 500 ft plots adjacent to streams. The pre-harvest and as-harvested results were measured in the field. The FMP and FPA harvests were simulated harvests of the pre-harvest data. All private sites are to the right of the orange line; state forest sites are on the left. Vertical lines indicated the difference between as-harvested and the potential harvest level at a site.

2.2.5 Harvest by distance from stream

The percent basal area reduction scenario demonstrated that a small fraction of basal area removed from within 100 ft of the stream would cause temperature increases above 0.3 °C. We examined how reduction in basal area as a function of distance from stream would affect temperature change. Due to differences in distance measure between state forests and private timberlands, the harvest was conducted according to horizontal distance (Figure 20) and slope distance (Figure 21). Slope distance is either equal to or greater than horizontal distance. Figure 21 includes slope distances out to 120 ft as this captures all but 1.5% of trees within 100 ft horizontal distance of the stream. The other trees had slope distances >120 ft at horizontal distances of 100 ft.

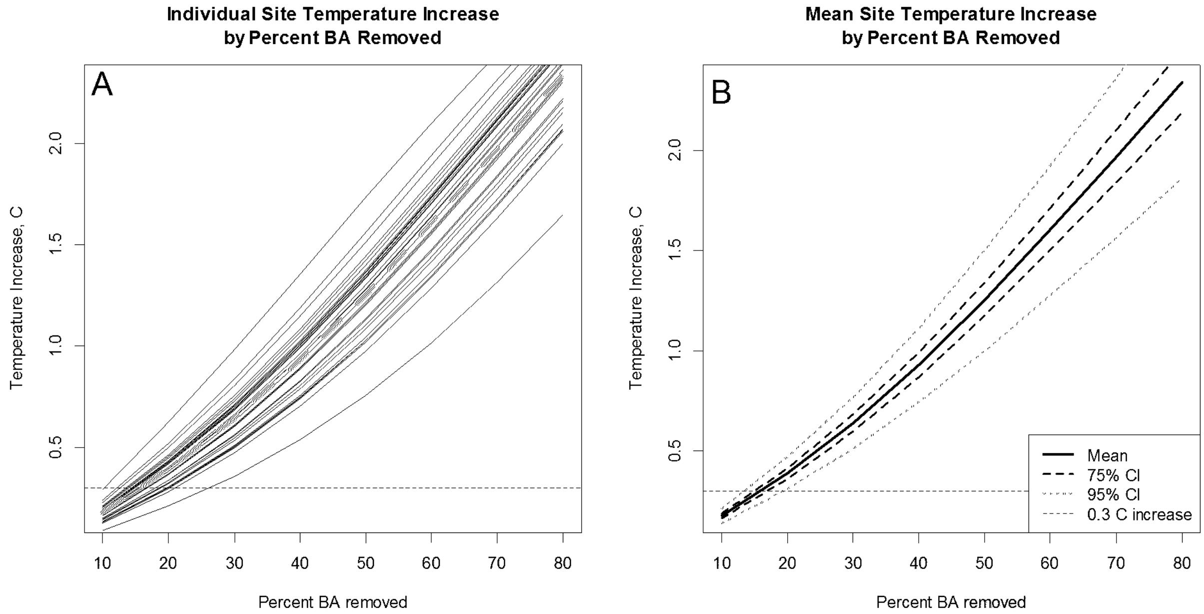

Figure 19. Temperature increases at a site level (A) and for an overall mean (B). The dashed horizontal line represents 0.3 ^{o}C. The orange line in A is equivalent to the line of the mean response in B. The blue and orange lines in B represent respectively the 75% and 95% Credibility Interval. The X axis for both graphs represents the percent basal area removed from each site. The Y axis is the temperature increase (^{o}C) due to the simulated harvests.

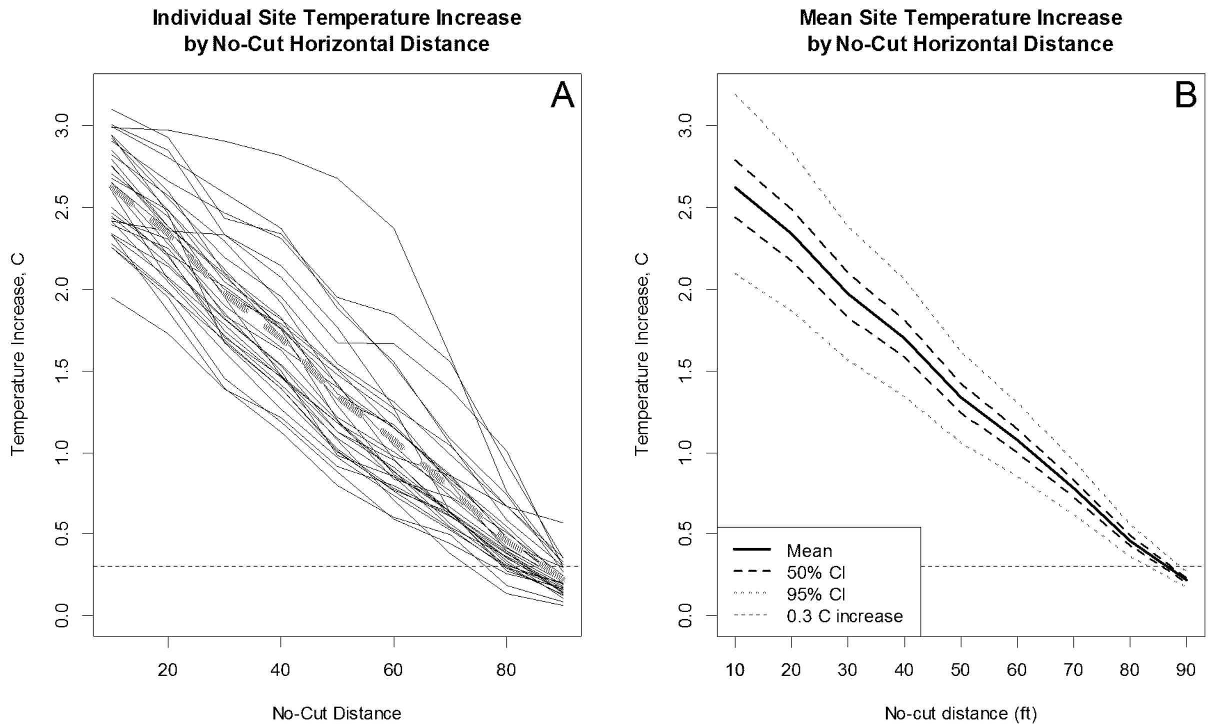

Figure 20. Predicted temperature increase at a site level (A) and overall mean (B) for harvest beyond specified horizontal distances. The orange line in A is equivalent to the line of the mean response in B. The blue and orange lines in B represent respectively the 75% and 95% Credibility Interval. The X axis for both graphs represents the no-cut distance that was not harvested for both banks of every site. The Y axis is the predicted temperature increase (^{o}C) due to the simulated harvests.

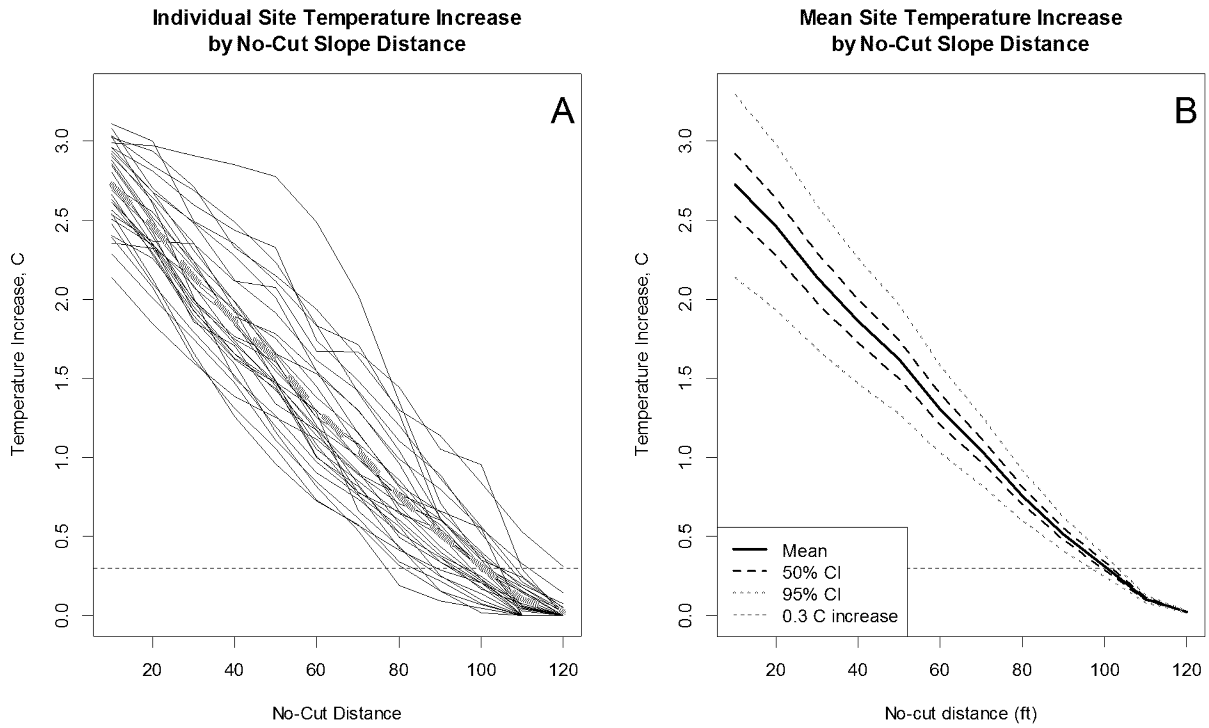

Figure 21. Predicted temperature increase at a site level (A) and overall mean (B) for harvest beyond specified slope distances. The orange line in A is equivalent to the line of the mean response in B. The blue and orange lines in B represent respectively the 75% and 95% Credibility Interval. The X axis for both graphs represents the no-cut distance that was not harvested for both banks of every site. The Y axis is the predicted temperature increase (^{o}C) due to the simulated harvests.

Literature Cited

Ellison, A. M. 1996. An introduction to Bayesian inference for ecological research and environmental decision-making. Ecological Applications 6:1036-1046.